WRDC-TR-89-1040

HIGH SPEED BUS TECHNOLOGY DEVELOPMENT

Marlan B. Modrow
Donald W. Hatfield

Rockwell International
Collins Government Avionics Division
400 Collins Rd. NE
Cedar Rapids, IA 52498

September 1989

Final Report for Period September 1983 - August 1988

AVIONICS LABORATORY
WRIGHT RESEARCH AND DEVELOPMENT CENTER
AIR FORCE SYSTEMS COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OHIO 45433-6543

NOTICE

When Government drawings, specifications, or other data are used for any purpose other than in connection with a definitely related Government procurement operation, the United States Government thereby incurs no responsibility nor any obligation whatsoever; and the fact that the government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data, is not to be regarded by implication or otherwise as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture use, or sell any patented invention that may in any way be related thereto.

This report has been reviewed by the Office of Public Affairs (ASD/PA) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

This technical report has been reviewed and is approved for publication.

PHILIP C. GOLDMAN, Project Engineer
Advanced Integration Group
Avionics Laboratory

LESTER McFAWN, Chief
System Integration Branch
Avionics Laboratory

FOR THE COMMANDER

CHARLES H. KRUEGER, Director
System Avionics Division
Avionics Laboratory

"If your address has changed, if you wish to be removed from our mailing list, or if the addressee is no longer employed by your organization please notify WRDC/AAAS *, W-PAFB, OH 45433 to help us maintain a current mailing list".*

Copies of this report should not be returned unless return is required by security considerations, contractual obligations, or notice on a specific document.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

1a. REPORT SECURITY CLASSIFICATION UNCLASSIFIED	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION/AVAILABILITY OF REPORT Approved For Public Release; Distribution Is Unlimited.
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S) WRDC-TR-89-1040
6a. NAME OF PERFORMING ORGANIZATION Rockwell International Collins Gov't Avionics Div 6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Avionics Laboratory Wright Research and Development Center
6c. ADDRESS (City, State, and ZIP Code) 400 Collins Rd. NE Cedar Rapids, IA 52498	7b. ADDRESS (City, State, and ZIP Code) WRDC/AAAS-1 Wright-Patterson AFB, OH 45433-6543
8a. NAME OF FUNDING/SPONSORING ORGANIZATION 8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER F33615-83-C-1036
8c. ADDRESS (City, State, and ZIP Code)	10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO. 63109F; PROJECT NO. 2734; TASK NO. 02; WORK UNIT ACCESSION NO. 03

11. TITLE (Include Security Classification)

High Speed Bus Technology Development

12. PERSONAL AUTHOR(S)

Marlan B. Modrow Donald W. Hatfield

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Final	FROM 9/83 TO 8/88	Sept 1989	162

16. SUPPLEMENTARY NOTATION

17. COSATI CODES		
FIELD	GROUP	SUB-GROUP
01	03	03
20	06	01

18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)

High Speed Data Bus, Avionics Architecture, PAVE PILLAR, Multiplex, ~~MIL-STD-1553~~, Fiber Optics, Local Area Network. (RH)

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

This report describes the development and demonstration of the High Speed Data Bus system, a 50 Million bits per second (Mbps) local data network intended for avionics applications in advanced military aircraft. The Advanced System Avionics (ASA)/PAVE PILLAR program provided the avionics architecture concept and basic requirements. Designs for wire and fiber optic media were produced and hardware demonstrations were performed. An efficient, robust token-passing protocol was developed and partially demonstrated.

The report covers the requirements specifications, the trade-offs made, and the resulting designs for both a coaxial wire media system and a fiber optics design. Also, the development of a message-oriented media access protocol is described, from requirements definition through analysis, simulation and experimentation. Finally, the testing and demonstrations conducted on the breadboard and brassboard hardware is presented.

20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☒ SAME AS RPT. ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED
22a. NAME OF RESPONSIBLE INDIVIDUAL Philip C. Goldman	22b. TELEPHONE (Include Area Code) (513) 255-4709 — 22c. OFFICE SYMBOL WRDC/AAAS-1

DD Form 1473, JUN 86 Previous editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE

UNCLASSIFIED

FOREWORD

This report describes the requirements generation, design, development, test, and demonstration of technologies associated with High Speed Data Bus (HSDB) applications for high performance military aircraft. The work was performed under Air Force contract F33615-83-C-1036 by Collins Government Avionics Division of Rockwell International. Fiber optic transmitter and receiver unit design and network characterization were performed by FiberCom, Inc. under subcontract to Rockwell. The principal engineering manager for Rockwell was Delaine Sather; Philip Goldman was the Air Force program manager. Merrill Ludvigson was the program manager at Rockwell; Kenneth Ferris performed the similar function at FiberCom. Kevin Milton and, later, Marlan Modrow acted as technical manager at Rockwell. Donald Hatfield at Rockwell and Philip Couch at FiberCom were the respective project engineers. Other Rockwell personnel who contributed much to the program include Ron Coffin, Mel Rhodes, Bob Jakoubek, Bob Wolter and John Senko.

TABLE OF CONTENTS

Paragraph Number	Description	Page
1.0	INTRODUCTION	1
1.1	Accomplishments	2
1.2	Program Overview	2
1.3	Report Organization	4
2.0	DESIGN OF A HSDB FOR USE ON AIRCRAFT	6
2.1	Determination of Candidate Approaches	7
2.2	Selection of the PAVE PILLAR HSDB Topology	10
2.2.1	Broadcast Bus vs. Ring Bus Trade Study	10
2.3	Selection of the PAVE PILLAR Protocol	13
2.3.1	Survey of Airframes	14
2.3.2	Synthesized HSDB Requirements	17
3.0	DEVELOPMENT OF COAXIAL NETWORK TECHNOLOGIES	18
3.1	Network System Design	18
3.1.1	Defining the Loss Budget	20
3.1.2	Determining Receiver Sensitivity Requirement	26
3.1.3	Determining Transmitter Power Output Requirements	29
3.1.4	Determining Receiver Dynamic Range Requirements	29
3.1.5	Network Operating Envelope	29
3.1.6	Modulation Characteristics	30
3.2	Coupler Development	31
3.3	Transmitter Design	34
3.4	Receiver Design	35
3.5	Proof of Concept Testing	38
4.0	DEVELOPMENT OF FIBER OPTIC NETWORK TECHNOLOGIES	47
4.1	Network System Design	48
4.1.1	Analysis Phase	49
4.1.1.1	Power Budget Analysis	50
4.1.1.2	Topologies Analysis	58
4.1.1.3	Optical Interconnect Analysis	67
4.1.2	Synthesis Phase	70
4.2	Transmitter Design	75
4.3	Receiver Design	77
4.4	Proof-of-Concept Testing	81
5.0	DEVELOPMENT OF THE PAVE PILLAR PROTOCOL	86
5.1	Requirements Survey	87
5.2	Baseline Protocol Design	89
5.3	Optimization of the Protocol	91
5.3.1	Message Data Base	101
5.3.2	Simulation Tool	103
5.4	Description of the PAVE PILLAR Protocol	105
5.4.1	Summary of the PAVE PILLAR HSDB Protocol	108
5.4.2	Evolution of the PAVE PILLAR HSDB Protocol	118
5.5	Verification/Validation of PAVE PILLAR HSDB Performance	119
5.5.1	Breadboard Terminal	119

(Preceding Blank)

5.5.2	Brassboard Terminal	120
5.5.2.1	Architecture	120
5.5.2.2	HSDB Card Design	123
5.5.2.3	Message Handler Card Design	126
5.5.2.4	TRU	128
6.0	TESTING, CHARACTERIZATION AND DEMONSTRATION	130
6.1	Validation of Coaxial HSDB Characteristics	132
6.2	Validation of Fiber Optic HSDB Characteristics	138
6.3	Validation of HSDB Protocol Characteristics	145
6.3.1	Message Transmission and Reception	146
6.3.2	Network Management	149
6.3.3	Service Functions	153
6.4	HSDB System Demonstration Equipment	155
7.0	RESULTS, CONSLUTIONS AND RECOMMENDATIONS	157
	REFERENCES	160
	BIBLIOGRAPHY	161

LIST OF FIGURES

Figure Number	Description	Page
1	Overview of the HSB Technology Development Program	3
2	HSDB Topology Candidates	9
3	Ringbus Network Produced by Rockwell in the Early 1970's	12
4	CSMA/CD Protocol Collapses as Offered Traffic Increases	16
5	Coaxial HSDB uses a Linear Bus Topology	19
6	Error Probability as a Function of Signal-To-Noise Ratio	20
7	Power Budget for the Coaxial HSDB Network	21
8	Loss Characteristic of Common Coaxial Cable Types	24
9	Measured vs. Specified Loss of RG-213 Coaxial Cable	25
10	Impact of Size, Impedance and Dielectric on Attenuation	25
11	Summary of S/N Analysis for a Typical Coaxial Network	27
12	EMI Induced on a Length of Coaxial Cable	28
13	Spectrum of MSK and Manchester Biphase Modulated Signals	30
14	Simplified Schematic Diagram of the Coaxial Network Coupler	32
15	Receive Port Coupler Equivalent Circuit	32
16	A Dual Diode Switch Disconnects this Transmitter Port from the Main Bus	33
17	Transmitter Port Equivalent Circuit in Transmit Mode	33
18	Transmitter Port Equivalent Circuit in Receive Mode	34
19	Lumped Constant Transmission Line Section	34
20	Prototype Wire Bus Coupler	35
21	The Coaxial HSDB Transmitter uses a Hybrid Circuit Power Amplifier	35
22	Manchester Waveform for '1001' Bit Pattern	36
23	Simulated Degration of a '1001' Bit Pattern at the End of the Network	36
24	Receiver Functional Block Diagram	37
25	Circuit Used for Ringing Tank Simulation	38
26a	Simulation Input Waveform	39
26b	Simulation Output Waveform	39
27	Breadboard Test Waveforms Validated Ringing Tank Clock Recovery Approach	40
28	Detection Circuit	40
29	Detector Performance on 'Near' Signals	41
30	Detector Performance on 'Far' Signals	41
31	Brassboard Coaxial HSDB Receiver Card	42
32	Brassboard Coaxial HSDB Transmitter Card	42
33	Brassboard Coaxial T/R Unit BER Characterization	44
34	HSDB System Demonstration Equipment Configured with Coaxial T/R Units and Couplers	45
35	Typical Signals Monitored on the Coaxial HSDB System Demonstration Equipment	46
36	The Fiber Optic HSDB Topology is a Single Transmissive Star	48
37	Summary of Fiber Optic System Constraints	52
38	Power Output Capabilities of Different Sources at 100 MHz	53
39	LED and Laser Output vs. Drive Current	53
40	Typical Laser Drive Compensation Circuit	54
41	An Avalanche Photodiode Requires Controlled Bias and Temperature	57

42 The Star Coupler Configuration is the Only Passive Candidate Meeting the System Loss Budget 60
43 Linear Bus Coupler Configurations 61
43a Typical Unidirectional Coupler 61
43b Typical Bidirectional Coupler 61
44 The Simplest Form of Linear Bus Topology Requires Bidirectional Couplers 62
45 Unidirectional Passive Linear Bus Candidate 62
46 Two Configurations of Optical Star Couplers 64
47 The Single Transmissive Star Topology Minimizes Loss 65
48 This Active Star-Star Topology Features Redundancy 65
49 Quad-Cluster Star-Star Topology Network Loss 66
50 The Star-Loop is a Hybrid of Linear Bus and Star Topologies 67
51 The Loop-Star Topology Supports Isolated Clusters of Nodes with Minimum Interconnect 67
52 This Aircraft Installation was Assumed for Determining the Network Loss Requirement 71
53 LED Coupled Power vs. Temperature 73
54 Fiber Optic Transmitter Simplified Schematic Diagram 76
55 Transmitter Output Power After Compensation 76
56 Three Different Approaches to Receiver Coupling 78
57 Fiber Optic HSDB Receiver Functional Block Diagram 78
58 A Transimpedance Preamplifier Provides Low-Noise Amplification of the Photocurrent 79
59 Brassboard Fiber Optic Receiver Circuit Card 81
60 Brassboard Fiber Optic Transmitter Circuit Card 82
61 Fiber Optic HSDB Network Emulator 83
62 HSDB System Demonstration Equipment Configured with Fiber Optic TRUs and Network Emulator 84
63 Fiber Optic Token Passing Network in Operation 85
64a Message Length Distribution 88
64b Message Repetition Rate Distribution 88
65 Distribution of Offered Traffic Short Term Rate 93
66 History of Short Term Message Delay 93
67 Message Delay vs. Offered Traffic Rate 94
68 Message Delay Characteristics of Several Protocol Overheads 95
69 Message Delay vs. Mean Message Length 96
70 Message Delay vs. Maximum Message Length 97
71 Message Delay vs. Token Rotation Timer Setting 99
72 Message Delay vs. Percentage Priority Traffic 101
73 Distribution of Average Message Length vs. Rate 102
74 HSDB Network Simulation Model Logic 104
75 Simulator Output Example 106
76 Percentage of Message Delay vs. Offered Traffic Rates 107
77 HSDB Network Generalized Topology 108
78 Packet Construct Requirements 109
79 Data Message Construct Requirements 110
80 Token Message Construct Requirements 111
81 Token Passing Information Transfer Format 112
82 Breadboard Terminal 120
83 Brassboard Terminal 121
84 Brassboard Terminal Major Components 122
85 HSDB Card Functional Block Diagram 124
86 Message Handler Functional Block Diagram 127

87	The TRU was Updated from the Task II Design to Make it More Producible	129
88	TRU Characterization Equipment	131
89	Coaxial Transmitter Output Waveform Requirement	133
90	Synchronization Waveforms are Illegal Manchester Symbols	134
91	Receiver Dynamic Range Test Setup	136
92	TRU Bit-Error Rate Test Setup	137
93	Bit-Error Rate Characterization Results	137
94	Fiber Optic Transmitter Waveform Requirement	139
95	Fiber Optic Synchronization Waveforms for Task II TRU	140
96	Fiber Optic Receiver Dynamic Range Characterization Setup	141
97	Fiber Optic Receiver Sensitivity vs. Temperature	142
98	E1 Brassboard Terminal Test Setup	144
99	Brassboard Semi-Custom Chip Set with Additional Support Chips	145
100	20-Word Message with Concatenated Token	147
101	The Concatenation Delimiter ED/SD Separates Messages	148
102	Brassboard Terminal Connected in the Data Generator Test Configuration	149
103	Three Node HSDB in Operation	150
104	Expanded Token Packet Showing Construction	151
105	The Terminal Response Time Requirement has been Validated	152
106	The Collision Vie Process has been Demonstrated	153
107	Demonstration of the Handshake between the Ringmaster and a Terminal Requesting Network Entry	154
108	Demonstration of Dropped Token Recovery by Retry of the Token Pass	155
109	The System Demonstration Equipment Block Diagram	156

LIST OF TABLES

Table Number	Description	Page
1	Alternative Topologies and Protocols	7
2	Candidate Topologies and Protocols	8
3	HSDB System Requirements	17
4	Summary of Coaxial Network Specification Requirements	22
5	Summary of Fiber Optic Network Specification Requirements	49
6	Technology Drivers	50
7	Fiber Optic Analysis Summary	51
8	Photodiode Characteristics	56
9	Topology Analysis Basic Loss Equations	59
10	Loss of Couplers	61
11	Network Loss Calculations for a Bidirectional Passive Bus	62
12	Network Loss Calculations for a Unidirectional Passive Bus	63
13	Quad Star-Star Network Loss Equations	66
14	Fiber Waveguide Tradeoffs	69
15	LED Qualification Test Results	73
16	Message Inventory Characteristics	87
17	PAVE PILLAR HSDB Protocol Summary	90
18	Iterative Process of Setting TRTs	100
19	Simulation Message Data Base	103
20	PAVE PILLAR vs. SAE AE-9B/L	108
21	Brassboard Semi-Custom Chip Set	146

LIST OF ABBREVIATIONS, ACRONYMS AND SYMBOLS

A/C	Aircraft
A/W	Amperes-per-Watt
AAMP	Advanced Architecture Microprocessor (Rockwell Proprietary)
AC	Alternating Current
AE-9	SAE Advanced Avionics Equipment and Integration Committee
AE-9B	SAE High Speed Data Bus Subcommittee
AE-9B/L	Linear Token Passing Bus Task Group
AF	Air Force
AGC	Automatic Gain Control
AIC	APM Input Controller
ALU	Arithmetic Logic Unit
APD	Avalanche Photodetector
APM	Access/Protocol Machine
AS-2	SAE Interconnect Networks Committee (Formerly AE-9B)
ASA	Advanced Systems Architecture
ASID	Advanced System Integration Demonstrations
ATA	Advanced Tactical Aircraft
ATF	Advanced Tactical Fighter
ATR	Acceptance Test Review
BDR	Bus Dynamic Range
BER	Bit Error Rate
BIT	Built-In Test
BIU	Bus Interface Unit
BW	Bandwidth
C	Celsius
CCITT	Comite Consultatif International Telegraphique et Telephonique
CDR	Critical Design Review
CDRL	Contract Data Requirements List
CMOS	Complementary Metal-Oxide Semiconductor
CPLR	Coupler
CRC	Cyclic Redundancy Check
CRT	Cathode Ray Tube
CSMA/CD	Carrier Sense Multiple Access/Collision Detect
DA	Destination Address
dB	Decibel
dBm	Decibels Relative to One Milliwatt
DC	Direct Current
DEM/VAL	Demonstration/Validation
DMA	Direct Memory Access
DPM	Data Processor Machine
DPRAM	Dual-Port Random-Access Memory
ECL	Emitter-Coupled Logic
ED	End Delimiter
EMI	Electromagnetic Interference
F.O.	Fiber Optic
FC	Frame Control
FCS	Frame Check Sequence
FET	Field-Effect Transistor
FIFO	First-In/First-Out
GFE	Government-Furnished Equipment
GHz	Gigahertz

HART	High Speed Data Bus Applications and Requirements Task Group
HDRM	Host Driver/Receiver Machine
HP	Hewlett-Packard
HSB	High Speed Bus
HSDB	High Speed Data Bus
Hz	Hertz (Cycles per Second)
IBM	International Business Machines
IC	Input Controller Chip
ICU	Input Controller Unit
IEEE	Institute of Electrical and Electronic Engineers
JIAWG	Joint Integrated Avionics Working Group
K	1024 (with Words or Bytes)
Kbps	Kilobits per Second
KHz	Kilohertz
km	Kilometer
L-C	Inductance-Capacitance
LAN	Local Area Network
LED	Light Emitting Diode
LHX	Light Helicopter-Experimental
LIT	Linear Implementation Task Group (SAE AE-9B/L)
μM	Micrometer
μS	Microsecond
mA	Milliampere
Mbps	Million Bits per Second (Megabits per Second)
MDS	Minimum Detectable Signal
MH	Message Handler
MHz	Megahertz
mM	Millimeter
mS	Millisecond
MSK	Minimum-Shift Keying
MUS	Minimum Usable Signal
mV	Millivolt
mW	Milliwatt
NA	Numerical Aperature
NEP	Noise-Equivalent Power
NIT	Network Inactivity Timer
NLT	Not Less Than
nm	Nanometer
NMT	No More Than
NRZ	Non-Return-to-Zero
nS	Nanosecond
OC	Output Controller Chip
OCU	Output Controller Unit
OSR	Optical Signal Range
PCM	Pulse-Code Modulation
PDR	Preliminary Design Review
pF	Picofarad
PI-Bus	Parallel Interconnect Bus
PIN	Positive-Intrinsic-Negative
PIO	Parallel Input/Output
PRE	Preamble
PROM	Programmable Read-Only Memory
RAM	Random-Access Memory

RCT	Reference Clock Time
RCV	Receiver
RCVR	Receiver
RF	Radio Frequency
RFI	Radio-Frequency Interference
RG-nnn	Military Part Number Designation for Coaxial/Triaxial Cable
RM	Redundancy Manager Chip
RMS	Root-Mean-Square
RMT	Ring Maintenance Timer
RMTM	Ringmaster/Topology Manager Chip
RMTMU	Ringmaster/Topology Manager Unit
RMU	Redundancy Manager Unit
ROM	Read-Only Memory
ROR	Receiver Operating Range
RTX	Receiver/Transmitter Chip
RTXM	Receiver/Transmitter Machine
RWT	Response Window Timer
Rx	Receive or Receiver
RXB	Receive Buffer
S/N	Signal-to-Noise
SA	Source Address
SAE	Society of Automotive Engineers
SCU	State Controller Unit
SD	Start Delimiter
SEM-E	Standard Electronic Module - Size E
SET	Solicit Entry Timer
SMA	Sub-Miniature A-Type Threaded Connector
SOW	Satement of Work
SPICE	Simuiation Program with Integrated Circuit Emphasis
SRR	System Requirements Review
TBD	To Be Determined
TNC	Threaded, Miniature Coaxial Connector
T/R	Transmitter/Receiver
TR	Technical Report
TRT	Token Rotation Timer
TRU	Transmitter/Receiver Unit
Tx	Transmit or Transmitter
Type N	Threaded, Standard Size Coaxial Connector
USAF	United States Air Force
V	Volt
Vp-p	Volts Peak-to-Peak
VSWR	Voltage Standing-Wave Ratio
W	Watt
WC	Word Count
WRDC	Wright Research and Development Center
XMT	Transmit
XMTR	Transmitter
Z_0	Characteristic Impedance

1.0 INTRODUCTION

The objective of the High Speed Bus Technology Development Program was to develop and demonstrate a high speed (50 Million bits per second) multiplex data bus for use on high performance military aircraft. This laboratory advanced development effort would hopefully lead to a military standard for a high speed bus.

With the introduction of many new electronics advances in the late 1970's, such as the microprocessor, the complexity of modern avionics systems has grown tremendously. Part of this trend has been the use of distributed processing. While this term has varied meanings, in this context it refers to the fact that computing power no longer must reside in a central unit because the costs have dropped so drastically. The result of the exploitation of this processing capability has been a radical change in the avionics architecture, or the way equipment and functions are grouped and partitioned. These new architectures require much more data exchange than would be the case with a centralized approach.

At about this same time, the Air Force had been enjoying the great success of a standard data exchange network or bus, MIL-STD-1553B. (1) By defining a standard, the possibility of a proliferation of solutions was eliminated and the ease of integration of equipment from different vendors was enhanced. Additional benefits of a data bus, as opposed to dedicated, point-to-point wiring, include improved flexibility, easier growth, reduced size, weight and power, and improved testability.

The High Speed Data Bus (HSDB) was planned to eventually take a place alongside MIL-STD-1553B as a mature, successful standard for avionics data distribution, with all of the benefits and significantly greater capabilities than 1553.

The USAF's Avionics Laboratory has long been a pioneer in avionics computer system standardization, both in hardware and software. It was instrumental in the development of MIL-STD-1553 along with MIL-STD-1589 (Jovial Language), and MIL-STD-1750 (Computer Instruction Set Architecture), and has been working to employ MIL-STD-1815 (Ada* Language). In the early 1980's the laboratory was involved in a major effort to define a next-generation avionics architecture employing distributed computing for better performance and improved fault tolerance. This effort was in two parts, the Advanced System Integration Demonstrations (ASID) program and later the Advanced System Avionics (ASA) program, nicknamed PAVE PILLAR. The High Speed Bus Technology Development program was initiated to support these efforts with the backbone communication network.

---

(1)MIL-STD-1553B, "Military Standard, Aircraft Internal Time Division Command/Response Multiplex Data Bus"

**Ada is a registered trademark of the Department of Defense (Ada Joint Program Office OUSR&E (RAAT).*

1.1 Accomplishments

At the time this program began, it was generally considered impossible to construct a reliable wire media linear bus with as many as 64 taps at a 50 Megabits per second (Mbps) data rate. By designing a bus coupler specially matched to the cable, this was successfully achieved. For the fiber optic media alternative, a state-of-the-art receiver was developed with excellent sensitivity and dynamic range performance. Finally, a bus access protocol was developed, in conjunction with industry and DoD experts working under the auspices of one of the Society of Automotive Engineers' (SAE) aerospace standardization committees, that is efficient, fault-tolerant, and which is suitable for a distributed avionics system.

As a result of this program, it has been shown that a 50 Mbps, 64-node HSDB network is feasible for next-generation production aircraft at minimal cost and risk.

Results from this contract have been continually provided to members of the SAE Committee AE-9B (later AS-2) who have been working to develop a HSDB standard. Significant work was performed on this program in the form of detailed hardware characterization and protocol simulation. Much of this was extremely valuable to the SAE committee. The HSDB designed under this program was detailed in a system specification, and is known as the PAVE PILLAR High Speed Data Bus.

The contractor teams now developing the Air Force's Advanced Tactical Fighter (ATF), the Navy's A-12 Advanced Tactical Aircraft (ATA) and the Army's Light Helicopter Experimental (LHX) have each developed specifications for a high speed data bus, each differing slightly from one another, PAVE PILLAR and the SAE's draft standard. Under the direction of a special group, the Joint Integrated Avionics Working Group (JIAWG), the contractors and the military services are working to develop standards across broad areas of avionics, with the HSDB being one. It is from this activity that a single HSDB standard will emerge.

1.2 Program Overview

The program was organized into four tasks. Tasks I, II and III were defined by the original contract with Rockwell. Task IV was added by contract modification when it became apparent that protocol technologies were tightly interrelated with the physical layer (transmitters, receivers and media) technologies.

Task I encompassed the development of transmitter/receiver units and couplers for a coaxial wire bus network. Task II included the development of transmitter/receiver units for a fiber optic bus network, using a star topology with a central optical star coupler. This task was largely performed by FiberCom, Inc. of Roanoke, VA under subcontract. Task III involved the

development of special test, characterization and interfacing equipment to support the other tasks. Task IV covered the development of the PAVE PILLAR protocol.

The technical program was accomplished over a 59-month period beginning with contract award on 30 September 83 and ending with the final demonstration review on 30 August 1988. Figure 1 summarizes the program chronology.

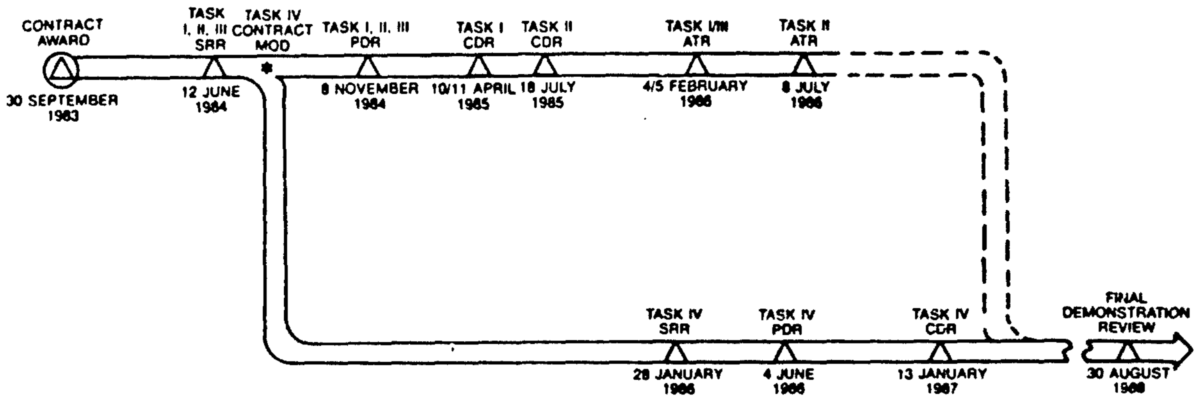

Figure 1. Overview of the HSB Technology Development Program

The original program plan was that Rockwell would focus on development of interconnect technologies and rely on the SAE High Speed Data Bus Committee to provide the protocol standard. This approach was initially supported by Rockwell in the form of a review of the SAE-generated HSDB requirements document and the preparation of a technical report, "Implementation Considerations for the High Speed Data Bus," distributed at their February 1984 committee meeting. When the SAE White Paper Evaluation Board selected a token-passing active ring approach for the HSDB, Rockwell was directed to stop work on a compliant approach and to begin independent effort to design a linear/star-based HSDB, including a full avionics-quality protocol. The addition of Task IV to the program by contract modification resulted from this directive. When the SAE later decided to support both active ring and passive linear bus working groups, Rockwell was directed to work with the SAE AE-9B/L (linear) task group. From this starting point Rockwell began development of transmitter/receiver hardware (Task I, II) and protocol. The working relationship between Rockwell and SAE remained supportive throughout the remainder of the contract, although the two designs diverged somewhat as it became apparent that the requirements for each were not identical. The driving force for Rockwell's effort was to create a highly-reliable design optimized for use aboard tactical aircraft; the SAE work attempted to solve a broad set of applications. Also, Rockwell's schedule forced decisions which in many cases were made prior to the similar decision on the part of the

SAE task group. In these instances Rockwell and the Air Force were reluctant to change due to the impact on already completed hardware and software. The result is that while the PAVE PILLAR HSDB and the SAE AE-9B/L HSDB offer similar performance and features, the two are not compatible and are not interoperable. This should not, however, lead to the assumption that one or the other effort was redundant. It is unlikely that either design would be as mature as is the case had it been developed singularly.

Task I and Task II work was completed in mid 1986 at completion of their respective acceptance test reviews (ATR), although the transmitter/receiver unit designs from Task II were used in the Task IV breadboard and brassboard terminal designs. Work on Task IV continued into late 1988 when the final demonstration review was held.

1.3 Report Organization

This final report is organized into several major topics as described below:

1.0 INTRODUCTION - An introduction and overview of the program, its goals and accomplishments.

2.0 DESIGN OF A HSDB FOR USE ON AIRCRAFT - This section describes the system level design studies carried out early in the program. These resulted in specification of the physical layer parameters, the broadcast topology, and the token passing protocol. This early systems design effort supported all four tasks.

3.0 DEVELOPMENT OF COAXIAL NETWORK TECHNOLOGIES - Describes the development of technologies associated with a HSDB network using coaxial cable media. This included the design of transmitter/receiver electronics and a bidirectional linear bus coupler. This work was accomplished under Task I.

4.0 DEVELOPMENT OF FIBER OPTIC NETWORK TECHNOLOGIES - Describes the development of technologies associated with a HSDB network using optical fiber media. This included the design of a state-of-the-art receiver and transmitter. This work was accomplished as a part of Task II.

5.0 DEVELOPMENT OF THE PAVE PILLAR PROTOCOL - Describes the development of the PAVE PILLAR HSDB protocol. The protocol was derived from the SAE AE-9B/L standard. Changes were made to provide better reliability in a high performance military aircraft. This work was accomplished as a part of Task IV.

6.0 TESTING, CHARACTERIZATION AND DEMONSTRATION - Describes the testing, characterization, and demonstration activities associated with the program.

7.0 RESULTS, CONCLUSIONS AND RECOMMENDATIONS - Self-Explanatory.

A brief summary of the program can be gained by reading the major section paragraphs (2.0, 3.0, etc.). Those readers interested in greater detail will find it organized into subordinate paragraphs in each section.

2.0 DESIGN OF A HSDB FOR USE ON AIRCRAFT

The Air Force has recognized a need for a high speed serial data bus for the next generation aircraft. Point-to-point wired interconnect weighs too much, occupies too great a volume and causes too many EMI/RFI compromises to be carried further. The first generation serial multiplex avionics bus, MIL-STD-1553B, solves many of these problems but offers only about 300 Kbps useful throughput. This forces the use of multiple data bus networks on present generation aircraft. Next generation systems must further improve upon installation flexibility, and must accommodate greater information flow, resource sharing, and fault tolerance. To satisfy that need the USAF awarded the High Speed Bus Technology Development contract to Rockwell International. The objective of this contract was to develop and validate HSDB network designs including T/R units, bus couplers, and media, and protocol operating at 50 Mbps. The designs were to be implemented in both wire and fiber optics.

The following basic requirements were assigned to the HSDB:

a. 50 Mbps information transfer rate
b. Up to 64 terminals
c. Up to 300 feet maximum terminal separation
d. 1 to 4096 work message length
e. 16 bit words
f. Point-to-point and broadcast modes

The following characteristics were not requirements, but were engineering goals:

a. Minimum message latency attributable to multiplexing
b. Compatible with both fiber optic and coaxial interconnect
c. Terminals added or deleted with no change to hardware or software of existing equipment
d. Active remote units minimized
e. 20 Mbps useful throughput in an operational system

Rockwell and FiberCom, under subcontract to Rockwell, completed various trade studies and analyses in order to arrive at the design described in the PAVE PILLAR HSDB system specification. It was determined that the PAVE PILLAR HSDB should be designed using broadcast topology and token passing protocol. These decisions are documented in subsequent paragraphs. The topologies and protocols considered will be described followed by a discussion of the application of both fiber optic and wire media to the selected topology and protocol.

2.1 Determination of Candidate Approaches

The analysis began with a survey of the leading alternatives for topology and protocol. The four alternative topologies and seven alternative protocols are shown in Table 1. Note that not all topologies are compatible with all protocols. This means that selection of a topology cannot be made independent of selection of the protocol. That point was not well understood at the time the original Statement of Work (SOW) for the program was prepared. A contract modification was ultimately processed to clarify the programs scope.

Table 1. Alternative Topologies and Protocols

TOPOLOGIES	PROTOCOLS						
	Command Response	CSMA/CD	Token Passing	Insertion Access	Time Slot	Request	Store & Forward
BROADCAST BUS	X	X	X		X		
FULLY CONNECTED	X	X	X		X		
RING	X		X	X	X		
SWITCHED NETWORK						X	X
"X" designates a compatible topology-protocol pair							

The thirteen topology-protocol pairs were analyzed and several were eliminated from further consideration based on performance deficiencies of an obvious nature.

- Store-and-forward protocols were eliminated on the basis of unacceptable message latency.
- Switched network topologies were eliminated on the basis of protocol complexity, the requirement for an active interconnect, and the difficulty of achieving broadcast message transfer.
- Fully connected topologies were eliminated on the basis of interconnect complexity and the difficulty of achieving broadcast message transfer.
- Insertion access protocols were eliminated because no reliable mechanism existed to control latency.

The remaining candidates are shown in Table 2. Note that most candidate protocols are compatible with each of the remaining topologies.

Table 2. Candidate Topologies and Protocols

TOPOLOGIES	PROTOCOLS			
	Command Response	CSMA/CD	Token Passing	Time Slot
BROADCAST BUS	X	X	X	X
RING BUS	X		X	X

Figure 2 illustrates the two remaining topologies. The broadcast bus is characterized as a network wherein a single node is transmitting at any time and each mode (including the transmitting node) hears the transmission in approximate real time. Note the broadcast bus may exhibit either a star interconnect or a linear interconnect. The ring bus is characterized as a series of point-to-point connections between successive nodes, which is closed on itself at the ends. Transmissions from the source node are repeated at each successive node around the ring until the message arrives back at the source node. Paragraph 2.2 describes the study which resulted in selection of the broadcast bus topology for the PAVE PILLAR HSDB.

Four candidate protocols remained for further study, as shown in Table 2. Command-Response protocol is characterized as having a single assigned node which grants permission to use the network according to some pre-determined algorithm. MIL-STD-1553B is an example of a command-response protocol. CSMA/CD (Carrier Sense Multiple access with Collision Detection) protocol is characterized as having each node transmit its message at the time it enters the queue, unless another node is already transmitting. Collisions are detected and transmission is retried in case of collision. IEEE 802.3 (2) (ETHERNET) is an example of a CSMA/CD protocol. Token passing protocol is characterized by each node waiting to transmit until it receives permission in the form of a special message (token) received from the node presently holding permission to transmit. The token circulates through the network according to an algorithm defined by the protocol. IEEE 802.4 (3) and SAE AE-9B/L Draft C (4) were used as examples of a token passing protocol. Time Slot protocol is similar to token passing protocol except that reception of the token is implied from waiting a defined period of time after a synchronization signal which is periodically received at all nodes. From these candidates token passing protocol was selected for the PAVE PILLAR HSDB. The study which resulted in this decision is described in paragraph 2.3.

---

(2)IEEE 802.3, "Carrier Sense, Multiple Access With Collision Detection"
(3)IEEE 802.4, "Token Passing Bus Access Method"
(4)SAE AE-9B/L, "Linear Token-Passing Multiplex Bus"

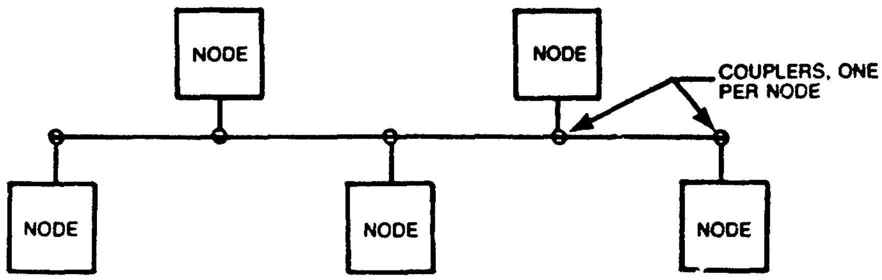

(a) BROADCAST BUS — LINEAR CONFIGURATION

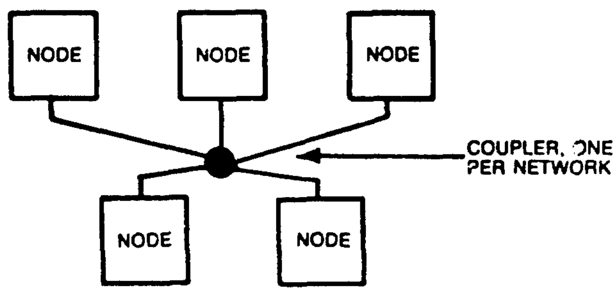

(b) BROADCAST BUS — STAR CONFIGURATION

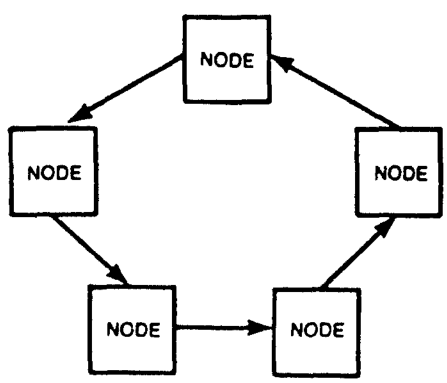

(c) RING BUS

Figure 2. HSDB Topology Candidates

2.2 Selection of the PAVE PILLAR HSDB Topology

The broadcast bus is preferred from the standpoint of reliability, installation, modification, and growth. The architecture is very simple and it is very easy to monitor activity on the bus. Good bus coupler design holds the key to a successful broadcast bus, because each terminal's transmitter must drive all other terminal receivers simultaneously and each terminal's receiver must deal with the accumulated loss, distortion, and reflection introduced by all of the other nodes. The active ring avoids these difficulties because it consists of a series of point-to-point connections. The ring approach, however, introduces its own set of potential problems such as:

a. Accumulated phase jitter
b. Fault location ambiguities
c. Logic errors are additive
d. Complexity of switching around idle or fault terminals

The PAVE PILLAR HSDB was implemented using broadcast topology because it better suited the requirements of an aircraft environment and because the problems and their solutions were well known, while the risk associated with potential problems and the cost for their resolution for the ring topology were deemed to be excessive.

2.2.1 Broadcast Bus vs Ring Bus Trade Study

When the contract was awarded, the Air Force intended to have Rockwell implement the high speed data base as defined by the SAE. Three topologies were still under consideration at that time. There was also concern about the upper data rate limit of the wire bus. Rockwell and FiberCom, under contract from Rockwell, prepared a white paper entitled, "Implementation Considerations for the High Speed Data Bus" which was distributed to the SAE White Paper Evaluation Board by the Air Force on 14 February 1984. In that report, Rockwell and FiberCom confirmed the feasibility of a 50 Mbps HSDB fiber optic bus utilizing LED optical sources and PIN detectors.

The SAE White Paper Evaluation Board met in Seattle in late February and evaluated four linear bus white papers and one ring bus white paper based upon the requirements defined by the High Speed Data Bus Application and Requirements Task Group (HART) in the "HART Requirements for the SAE AE-9B/L High Speed Data Bus" document. (5) As a result of those deliberations, the SAE White Paper Evaluation Board recommended that the token passing active ring bus configuration be adopted by the full HSDB subcommittee. As a result of that

(5)High Speed Data Bus Application And Requirements Task (HART) For The SAE AE-9B/L High Speed Data Bus

recommendation, Rockwell was directed to begin technology development for a token passing ring and prepare a plan for the presentation to the Air Force in early April.

During the report of the SAE evaluation board to the full HSDB subcommittee the board described the basis for its decision. The comparisons based upon "firm" requirements did not enter into the final recommendations. While there were differences in performance between the linear and the ring buses which would slightly favor the ring, they were not deemed sufficiently significant to base a decision, so "firm" requirements were put aside and both topologies were evaluated looking at the "desirable" characteristics defined in the HART document, using the "fuzzy decision making" process. The board described the process and the numerical matrices that resulted in its final decision but did not describe the bus configurations evaluated in detail nor the considerations that led to the numbers in the matrix, even upon request.

On 5 April 1984, a presentation was made by Rockwell to the Air Force which contained the following key points:

a. Rockwell has considerable prior active ring experience having designed, developed and delivered approximately 80 active ring message switch systems, operating at 32 Mbps with wire media, using biphase modulation. (Figure 3 shows a typical installation.)

b. While there was no doubt that we could design an active ring bus that would satisfy the Air Force needs there could be no assurance that when it was completed it would be as cost effective as a linear bus.

c. To bring the design level of maturity for active rings to that for linear buses requires that a number of trade studies be conducted concerning such issues as:

   - *Survivability:* This included ambiguities caused by decentralized control; terminal by-pass implementation; error recovery including analysis of the ambiguities of fault location, methods for redundancy, the potential for partitioning of the ring, the by-pass dynamic range problem, and the potential need for alternate routing and physical bus separation needs.
   - *Terminal characteristics:* This included issues having to do with changing bits on the fly, buffer size, terminal delay, elastic buffers in individual terminals and fixed-terminal delay with one elastic buffer, normal message handling, and error recovery.
   - Installation, modification and growth issues had to do with systematic wiring strategy and alternative bus separation.

4. These trade studies would require additional funds.

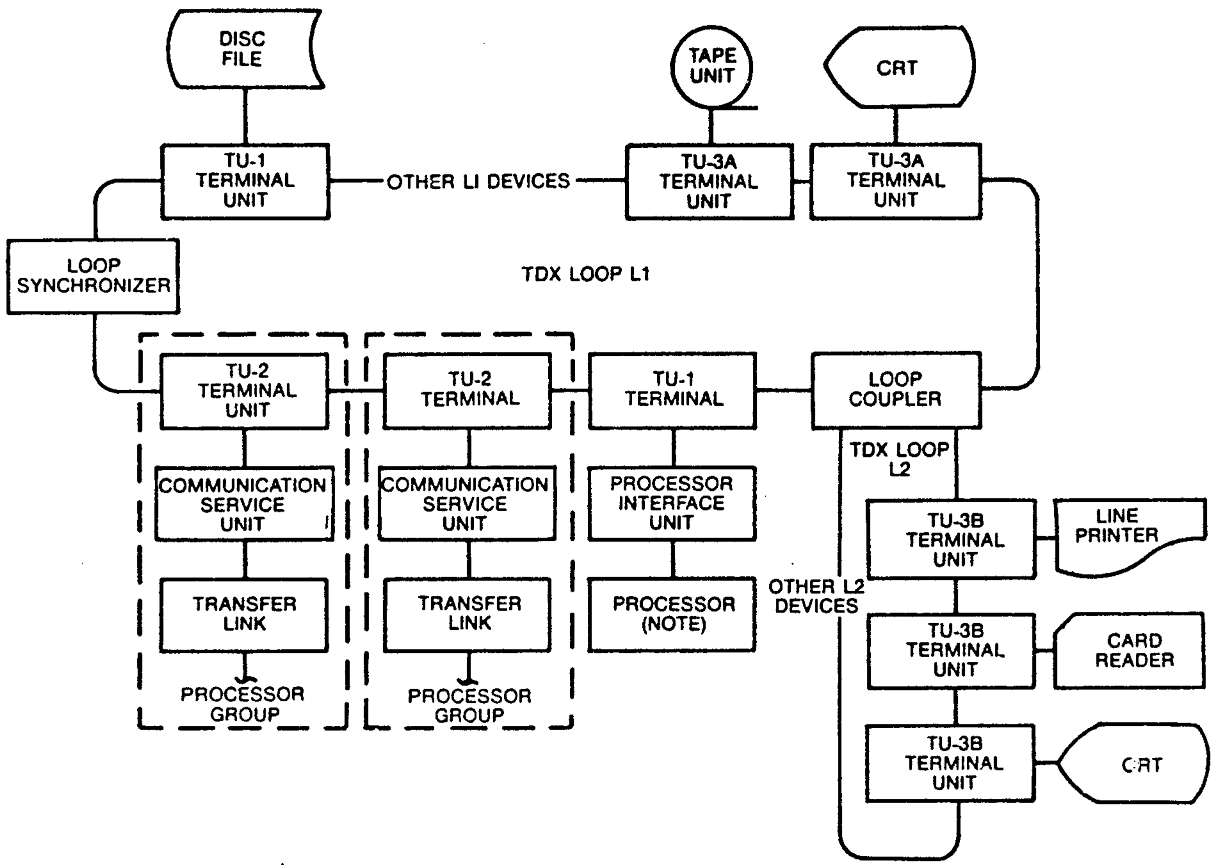

Figure 3. Ringbus Network Produced by Rockwell in the Early 1970's

Since additional funds were not available to complete the TR unit development effort described in our active ring plan, we were requested to prepare an alternative plan showing what could be accomplished with available funds. On 6 April 1984, we proposed a program which consisted of analysis and demonstration elements. The analysis phase would address six areas:

a. 128 terminals
b. Survivability estimates
c. Error rate budgets
d. Protocol/error recovery defined
e. Computer simulation
f. Installation/maintenance plan

Several areas of deficiency existed in the plan. Specifically the analysis would leave individual terminals uncharacterized and unproven. Clock recovery/tracking would remain undemonstrated as would by-pass characteristics. In summary, the analysis would carry limited credibility.

The following hardware demonstration elements were proposed:

a. 4 Ring members integrated with 4 MIL-STD-1750A processors
b. 4 By-pass units/couplers
c. Operation at 50 Mbps, fiber optic
d. Simple token passing protocol (contractor selection)

The following shortcomings were identified in the proposed demonstration:

a. 128 terminals not proven
b. Survivability unknown
c. Minimum error recovery
d. No budget for error rates
e. No demonstration of maximum by-pass
f. No A/C installation/maintenance/topology strategy
g. No necessary SAE AE-9B/L protocol
h. No round trip latency demonstration

Based upon the Air Force's assessment of the overall relative risks, decision was made to proceed with the linear bus. It should be noted that the recommendations of the SAE White Paper Evaluation Board were not approved by the vote of the full committee and as a result the SAE initiated plans to define both linear and ring bus configurations. All Rockwell interaction with the SAE from that point forward was with the Linear Task Group (LIT).

2.3 Selection of the PAVE PILLAR Protocol

At the time of contract award, the Air Force intended to have Rockwell implement HSDB protocol as defined by the SAE. This approach, in fact, was used to arrive at the choice of a token passing protocol. The details of the protocol were, however optimized for the PAVE PILLAR application. (This tailoring effort is described later in this report.) Three different protocols were being seriously considered at that time, token passing, command/response, and CSMA/CD.

Token passing protocols are very attractive because new terminals (subsystems) may be added to the network with no modification to existing terminals. Token passing (decentralized

control) and Command/Response (centralized control) are generically similar in that no one may use the bus without being passed a "free token". The difference is that in token passing, the "free token" is passed in a "logical ring" from one user to another, while in Command/Response the "free token" is passed only by a central controller terminal. Token passing is most efficient when the traffic is evenly distributed throughout the network. Efficiency decreases if many terminals are idle at any one time because they continue to handle tokens even if they have nothing to transmit. An attractive feature of token passing protocols is that terminals can be easily added to, or removed from, an active network without requiring any change to existing terminals.

The Command/Response protocol is easily understood. Its attractive feature is that a predesignated control node is always in charge. It is most efficient if many terminals are idle at any one time. The controller doesn't bother to interrogate those terminals until the control program determines that they will have something to transmit. The unattractive feature of command/response is that the bus controller software must be changed every time anything in the system changes. Also, it decreases in efficiency when traffic is evenly distributed. If most terminals will have activity regularly, then a lot of bus time is wasted by returning control to the bus controller between each message.

The CSMA/CD protocol is attractive because terminals (subsystems) may be added or deleted from the bus with no changes to any of the non-interfacing terminals. CSMA/CD works most efficiently when messages are long with respect to the media propagation time but deteriorates rapidly under heavy traffic loads because of the increase in retransmission traffic, and ultimately collapses.

The PAVE PILLAR HSDB using a token passing protocol, operates efficiently under the type of traffic load conditions expected aboard an aircraft. The design also allows subsystems to be logically decoupled from one another which minimizes the cost and risk of aircraft modifications.

2.3.1 Survey of Airframes

Final definition of the PAVE PILLAR HSDB protocol resulted from optimization of the SAE AE-9B/L protocol for the PAVE PILLAR application. The process of optimization proved to be quite involved. The reason for this was the low level of experience with local area network (LAN) technology aboard aircraft. No reliable data set existed for the application because no present generation aircraft uses a LAN in a manner similar to that defined by the PAVE PILLAR architecture. In the absence of hard data in this area, Rockwell used a survey of the seven Advanced Systems Avionics (ASA) contractors, augmented with data from other appropriate sources, to down-select to a single protocol type. Several topics were of special interest during the survey:

a. The level of distribution of control functions throughout the aircraft. A single dedicated mission processor defined one endpoint. A fully distributed network where each node performed flexibly as one part of the mission control function defined the other endpoint.

b. The need for service functions; these might include global reference clock, guaranteed delivery, automatic notification of error, crypto, functional and broadcast addressing modes, automatic configuration of processors, automatic time tagging of messages, adaptive network tailoring, and other similar functions not directly associated with communication between two users.

c. The number of nodes active at any time, the number of messages initiated per unit time, size of messages, required latency and other performance related characteristics.

d. Reliability requirements for individual messages and for the network itself; the impact of undetected errors, and the maximum recovery time for dead network.

e. Installation issues including partially populated systems, growth, and different avionics configurations.

f. The level of interface with the user. Should the HSDB be a 'smart' network where the user merely deposits a message in the mailbox or should it be highly interactive with the user process?

The survey yielded an unexpected result. Much of the commercial work done on LAN design was found to be not applicable for the PAVE PILLAR application. The principal tenant of LAN design is very different in the two environments: (1) Commercial LAN designers assume they have little or no control over applications; (2) Aircraft LAN designers insist on much tighter control over the system. This difference in philosophy shows up in a variety of different ways but one which is easy to describe and grasp is that of protocol adaptation as the network operates. Commercial designers must assume that application after application will be added to the network and that each application is 'selfish' in that it gives highest priorities to its own network access, lower to any other application. The Commercial LAN designer must provide protocol features which act to distribute network access among all users in an equitable manner no matter what priority any specific application requests. This implies the ability to closely monitor network activity and to provide real-time adaptations to fine-tune the network as operating conditions change.

Aircraft LAN systems designers on the other hand, insist on maintaining very tight control over each and every application using the network. A network global priority definition results from this desire to maintain a deterministic environment. As a result, network monitoring

is of little importance except as a development tool and real time adaptations are not needed, or even desired. This may eventually change but for next generation aircraft, there is definitely a philosophy of very tight management of network operation.

Also, the critical nature of the application, tended to neutralize the significant amount of work completed in the commercial/industrial sector on analysis of protocols. For example, most commercial protocol characterization has used average latency as a fundamental trade study characteristic. Airframes, on the other hand, were little interested in average latencies. They emphasized worst-cost performance and wished to establish guaranteed delivery times. Network crashes, while a nuisance in commercial networks are treated as one element of a cost-benefit trade and allowed to happen at some rate in order to optimize other characteristics of the network. On an aircraft, conversely, a network crash, even of short duration, could cause loss of the platform, and possibly loss of life.

As a result of the survey the choice of a token passing protocol was revisited and affirmed. This came about more as the result of deficiencies in the other two candidates than an obvious superiority of token passing.

CSMA/CD was eliminated first since it exhibits an unacceptable discontinuity in performance under conditions of increasing load. This characteristic has been identified in simulations performed by Rockwell and others. Figure 4 illustrates this point by showing collapse of the network (infinite message delay) at a certain critical message arrival rate which depends

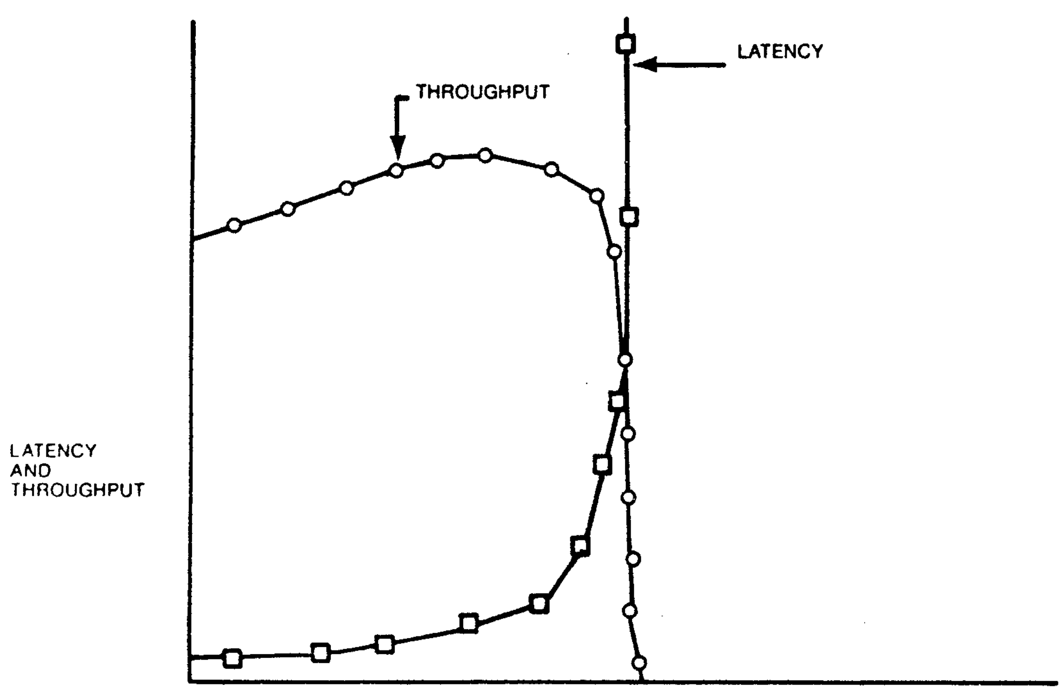

Figure 4. CSMA/CD Protocol Collapses as Offered Traffic Increases

upon message characteristics. This happens because collisions and retrys become an increasing message load as offered traffic increases. As a result, throughput decreases as offered traffic increases above the critical point, and eventually 100% of messages result in collision.

Command/Response protocols were also eliminated. Designing a 50 Mbps version of MIL-STD-1553B of the protocol is technically practical, but this candidate was not selected because it operates with significantly poorer efficiency in the type of architectures envisioned for PAVE PILLAR, i.e., many nodes with processing power and widely distributed control structure. This decision was not universally accepted by the ASA contractors, at first, but as simulation results showed the advantages of the token passing protocol opposition faded to the point where it was not a serious point of discussion towards the end of the program.

2.3.2 Synthesized HSDB Requirements

Table 3 lists the HSDB network requirements which resulted from the initial design exercise described in this section. These requirements governed follow-on development of coaxial network technology, fiber optic network technology, and protocol. The details of the development of each enabling technology is provided in Section 3 of this report.

Table 3. HSDB System Requirements

CHARACTERISTIC	REQUIREMENT
a. Data Rate	50 Mbps
b. Information Rate	20 Mbps in typical application
c. Number of nodes	2 through 64
d. Physical separation	100 meters maximum
e. Message length	4096 words maximum
f. Latency control	Message-by-message priority system with 20mS for highest priority messages
g. Addressing	Fixed, logical, broadcast
h. Interconnect	Broadcast, fiber optic with coaxial optional
i. Growth	Added or removed nodes with no change to operational nodes
j. Reliability	Less than 1 detected error per 400 seconds; less than 1 undetected error per 100 minutes
k. Global reference clock	Accuracy 1 part in 10^5
l. Protocol	Token passing
m. User interface	Simple asynchronous 'mailbox' interface (PI-Bus as the target design)

3.0 DEVELOPMENT OF COAXIAL NETWORK TECHNOLOGIES

Development of the coaxial version of the HSDB encompassed investigation of the technical limitations of the technology, establishing performance goals for the network, and design/fabrication/test of prototype hardware as proof of concept. This was performed as Task I of the HSB Technology Development Program and fulfills the requirements of paragraph 4.1 of the SOW. The linear bus topology was the preferred approach from the standpoint of installation and growth flexibility. There is no doubt that a star configuration is technically feasible but the linear bus configuration was selected, pending discovery of insurmountable technical problems, for the above mentioned reason.

Development of a linear coaxial network represented a difficult design problem. Whereas most other topologies take the form of point-to-point circuits between terminals, a linear bus, to satisfy all of the program requirements, must interconnect 64 terminals simultaneously through a passive network. As shown in Figure 5, the network consists of transmitters, receivers, couplers and coaxial interconnect. System loss and dynamic range are the principal characteristics limiting application of coaxial linear bus topologies. Secondary concerns involved determining the affect on performance of reflections and other distortion components. Early during the program it became apparent that achieving a superior coupler design was key to the success of the coaxial linear bus approach. Engineering work focused on development of a coupler with very low excess loss and well controlled impedance and group delay characteristics. System requirements forced development of the state-of-the-art designs which are described in paragraph 3.2. Other elements of Task I required the application of standard engineering techniques to arrive at a functional design.

3.1 Network System Design

The design for the coaxial HSDB network evolved through the multi step analysis described below:

1. A loss budget was established using technological limitations as a basis. The budget included cable loss, coupling loss, coupler mainline loss, and connector loss. Figure 5 illustrates the power budget arrived at by this method. Paragraph 3.1.1 describes this part of the analysis.
2. Receiver sensitivity was determined. First, the signal-to-noise (S/N) ratio necessary to meet system detected error rate specifications was established. This required definition of the sources and magnitudes of noise in the system. Figure 6 shows the design point. Paragraph 3.1.2 describes this part of the analysis.

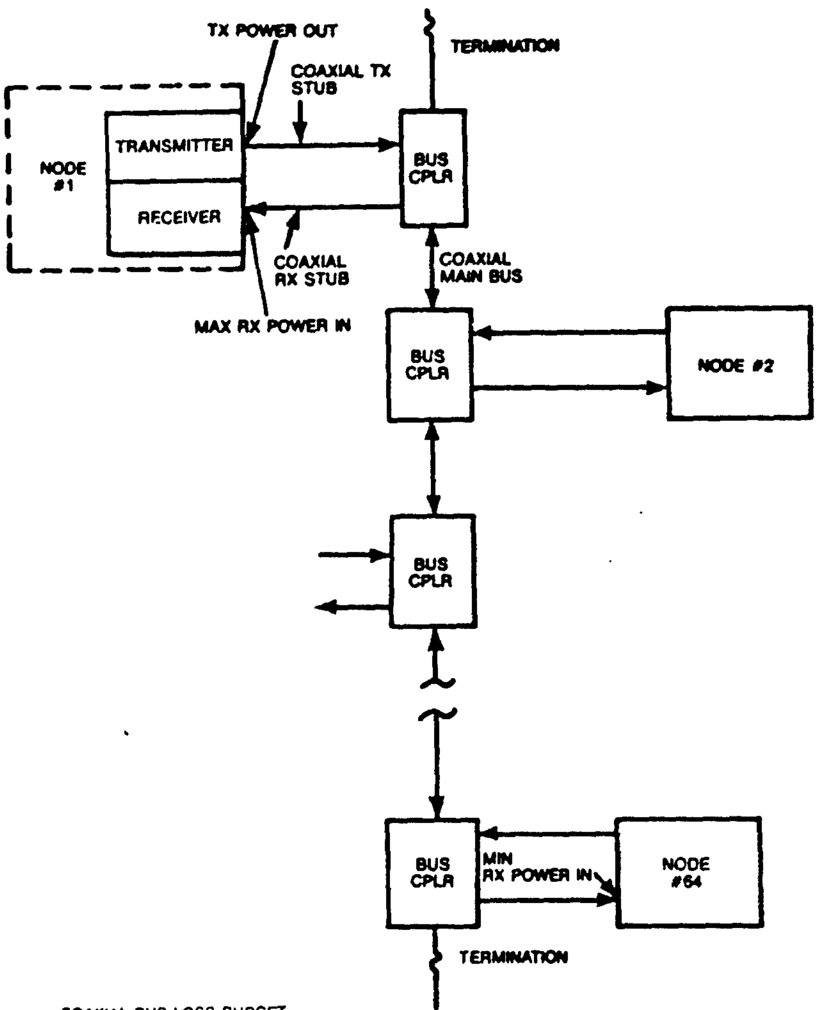

COAXIAL BUS LOSS BUDGET

TRANSMIT STUB AND COUPLING	10	dB	
RECEIVE STUB AND COUPLING	26	dB	
	36	dB	MINIMUM LOSS
—PLUS—			
COAX TRUNK LOSS	8	dB	
COUPLER LOSSES (64 × 0.2 dB)	12.8	dB	
CONNECTOR LOSSES (126 × 0.02 dB)	2.5	dB	
	59.3	dB	MAXIMUM LOSS

Figure 5. Coaxial HSDB uses a Linear Bus Topology

3. The necessary power from the transmitter was determined from worst-case signal power at the receiver and the worst-case network loss. Figure 7 shows the network power budget.
4. Receiver dynamic range was determined from the transmitter output specification, coupler specification, and network loss specification. Figure 7 also shows the receiver operating range and dynamic range derivation.

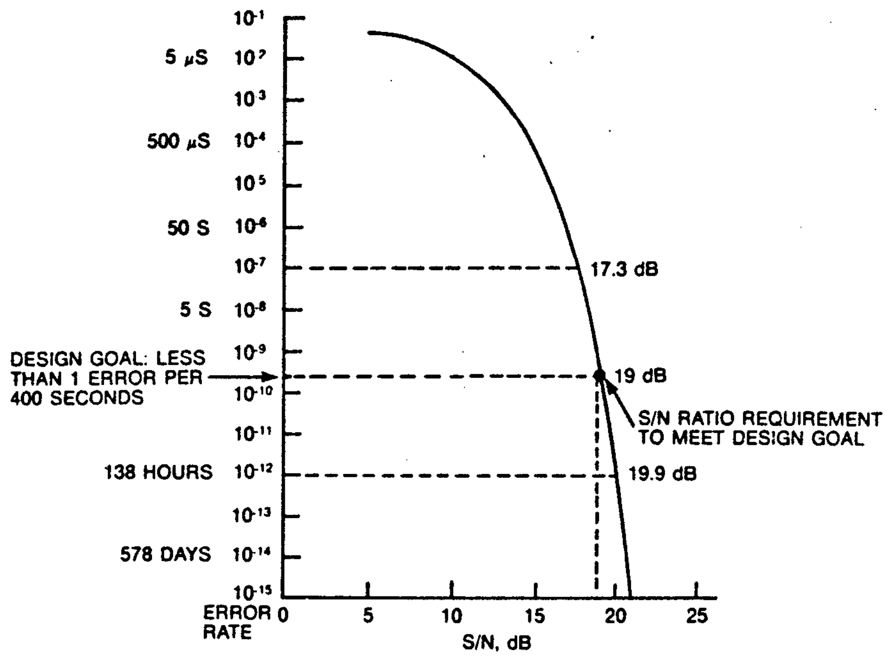

Figure 6. Error Probability as a Function of Signal-To-Noise Ratio (6)

As a result of this analysis, development specifications for the transmitter, receiver, and coupler were prepared, and development was initiated. Table 4 and Figure 7 summarize the resultant requirements. These specifications were somewhat modified during the course of the program as better technical understanding was achieved. Final values are found in the PAVE PILLAR HSDB system specification.

3.1.1 Defining the Loss Budget

Loss budget refers to the attenuation characteristic of the network between the transmitter output port and the receiver input port. The loss budget design goal for Task I is comprised of the elements defined below, and illustrated by Figure 5.

a. Transmitter stub loss (coax cable)
b. Tx port coupling (coupler)
c. Mainbus loss (coax cable plus couplers plus connectors)
d. Rx port coupling (coupler)
e. Receiver stub loss (coax cable)

---

(6)"PCM And Digital Transmission Systems," McGraw-Hill Company, 1982

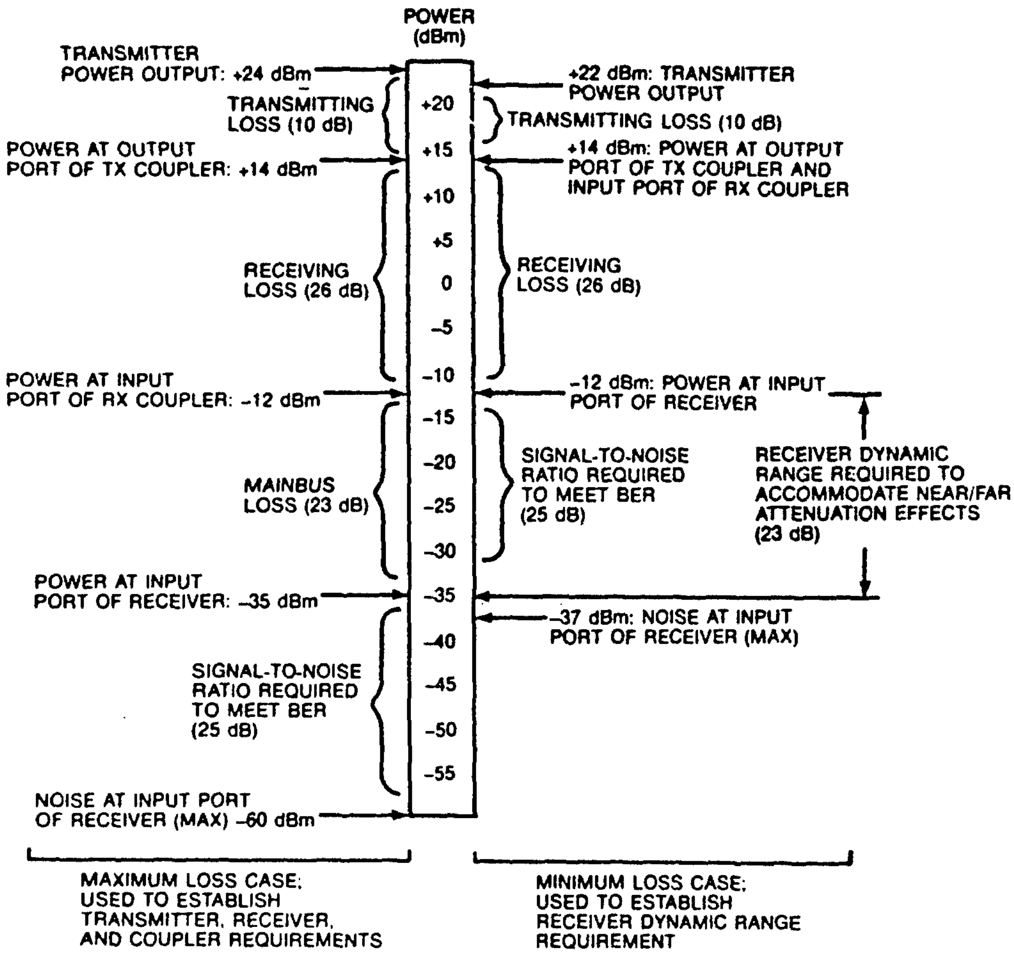

Figure 7. Power Budget for the Coaxial HSDB Network

Determining what values to assign to each element of the network was an essential first step of the network design task. The limitations of technology are well known for coax cable and for connectors. This meant that most parts of the analysis were quite straight forward. The coupler, on the other hand, represented a component for which no performance precedent existed. For this reason the final design goals were arrived at using an iterative process. The result of this process is described below.

Coupler

Three characteristics of the coupler were of importance for determining the system loss budget:

a. Coupling from the Tx port to the mainbus ports.

b. Coupling from the mainbus ports to the Rx port.

c. Loss from mainbus port to mainbus port.

Table 4. Summary of Coaxial Network Specification Requirements

Receiver Operating Range: The receiver shall meet all performance requirements while receiving transmissions whose 50 MHz component signal level is greater than 3.14 mV rms (-35 dBm) and less than 44.5 mV rms (-14 dBm).

Receiver Dynamic Range: The receiver shall meet all performance while receiving transmissions wherein packets are preceded by TBD offtime and 8 bits preamble, and alternate packets are at opposite receiver operating range points.

Transmitter Output: Signal out of the transmitter shall be 7.1 V rms (+24 dBm) ±TBD.

Coupler Tx Port Coupling: Coupling from the Tx port of the coupler to either mainbus port shall be 10±TBD dB.

Coupler Rx Port Coupling: Coupling from either mainbus port of the coupler to the Rx port shall be 26±TBD dB.

Coupler Insertion Loss: Loss through the mainbus of the coupler shall be less than 0.2 dB at 50 MHz.

Coupler Mainbus VSWR: Return loss of either mainbus port of the coupler shall be greater than 35 dB while the other mainbus port is terminated in 50 Ohms.

Coupler Mainbus Group Delay: The difference in delay time through the coupler shall be less than TBD nS between 25 MHz and 75 MHz.

Modulation Format: Manchester II

Characteristics (a) and (b) are principal design points. Any value, within reason, can be achieved by adjustments to the design. Characteristic (c) is comprised of two factor: coupling loss, which is the power on the mainbus less the power coupled into the Rx port; and excess loss, which is power lost within the coupler for various reasons.

The Tx port coupling requirement was eventually established as 10 dB. This value represents a compromise between two conflicting goals; (1) the desire to minimize the power required from the transmitter dictated selection of a low coupling constant and (2) the desire to isolate faults within the transmitter or transmit stub cable dictated choice of a high coupling constant. Ten dB represents the minimum coupling value which will not result in unacceptable degradation to other network nodes should a short occur in a transmitter or the transmit stub. The requirement for an electronic switch in each coupler to disconnect the transmit stub from the network while in receive mode of operation also resulted from this phase of the design task. The

switch provides two functions. First, it disconnects a shorted transmitter output or transmit stub from the coupler. Second, it eliminates any noise generated by the transmitter output stage from appearing on the mainbus thereby improving the S/N ratio on the mainbus.

Rx port coupling was eventually established as 26 dB in a similar manner. A low coupling constant is desirable because it minimizes the amount of transmitter power required. A high coupling constant is desirable for reliability and also because it minimizes the dynamic range required of the receivers. Selection of 26 dB represented a practical design point.

Loss on the mainbus coupler arm was specified as 0.2 dB. This includes both the loss from coupling power at -26 dB to the receive port and the excess loss of the coupler. Whether this could be achieved for prototype couplers was of great concern since no coupler of similar design had ever been built prior to this program. Calculation showed it to be within reason, however, so this became the target specification. Subsequent fabrication of prototype couplers showed this to be an attainable specification.

Coaxial Cable

As shown in Figure 5, loss of the coaxial cable interconnect represents the second largest component of the loss budget, 8 dB. This figure was budgeted from review of the characteristics of many different types of coaxial cable. A single type was not selected for use on the network because each application will have slightly different requirements. Rockwell felt that the system designer should be able to select an appropriate cable depending on the interconnect length, number of nodes, EMI/RFI requirements, temperature, etc.

Five characteristics of cable are of prime interest to the system designer when trying to optimize a system architecture. Those characteristics are:

a. Attenuation
b. Phase distortion or jitter
c. Velocity of propagation
d. Shielding effectiveness
e. Physical size

The attenuation of a number of available cables is shown in Figure 8. Measurements made on a 175 foot length of RG-213 cable illustrated that the specification was very conservative (see Figure 9) when compared with measured performance of actual cable. Figure 10 shows the impact that size, characteristic impedance, and the type of dielectric has on cable loss. From this data, it is obvious that for minimum loss the cable should be 75 Ohm cable with foam dielectric. Our original draft specification was for a 75 Ohm system. This decision was modified later, however, because of the poor availability of 75 Ohm cables qualified for use aboard aircraft. The

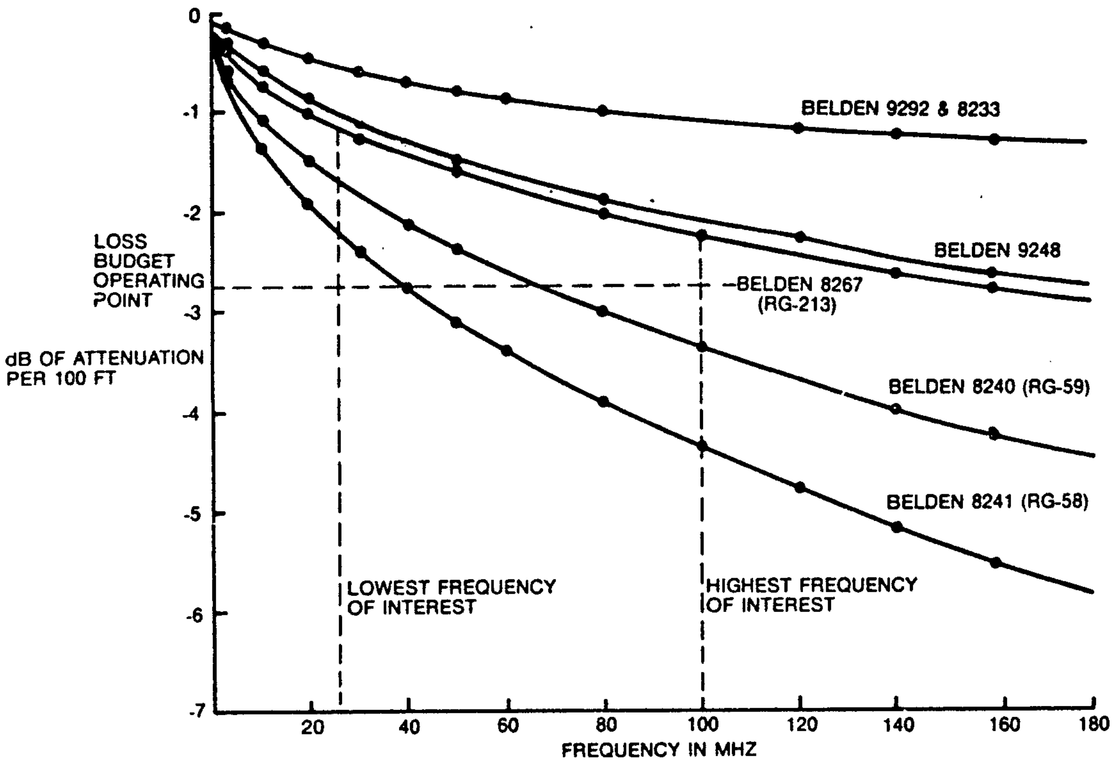

Figure 8. Loss Characteristic of Common Coaxial Cable Types

final design specified a 50 Ohm system because of the wide availability and experience with 50 Ohm coaxial cable in aircraft, and because analysis showed that all network requirements could be met using a 50 Ohm design.

Measurements and calculations were made to characterize the impact of the coax cable on performance of the network. These showed that the phase shift is not a problem. It is essentially constant over the frequency range of interest. The velocity of propagation does impact, to a small degree, the propagation delays, but is not a problem. Foam dielectric, such as foamed polyethylene, results in a propagation constant of 0.78, somewhat better than the 0.66 typical for solid polyethylene, but either could be accommodated. The differential loss across the bandwidth of interest was the major concern. This characteristic is shown in Figure 8. In a coaxial network the higher frequency components of the signalling waveform will be attenuated much more than the lower frequency components of the waveform. This differential loss variation caused additional complexities with the network design in two ways:

1. The S/N ratio design point was forced to the highest frequency component of the modulation envelope. This meant that the S/N ratio of other modulation components was in fact higher than specification. This caused testing difficulties.

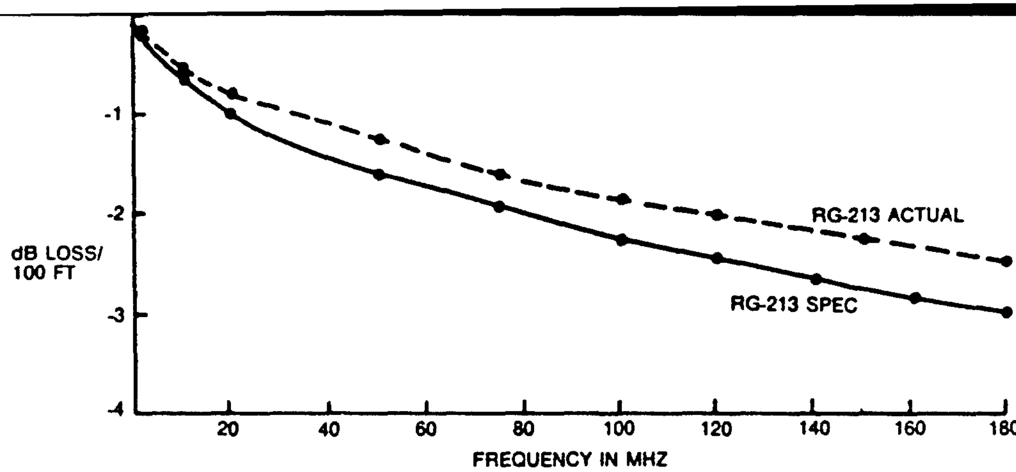

Figure 9. Measured vs. Specified Loss of RG-213 Coaxial Cable

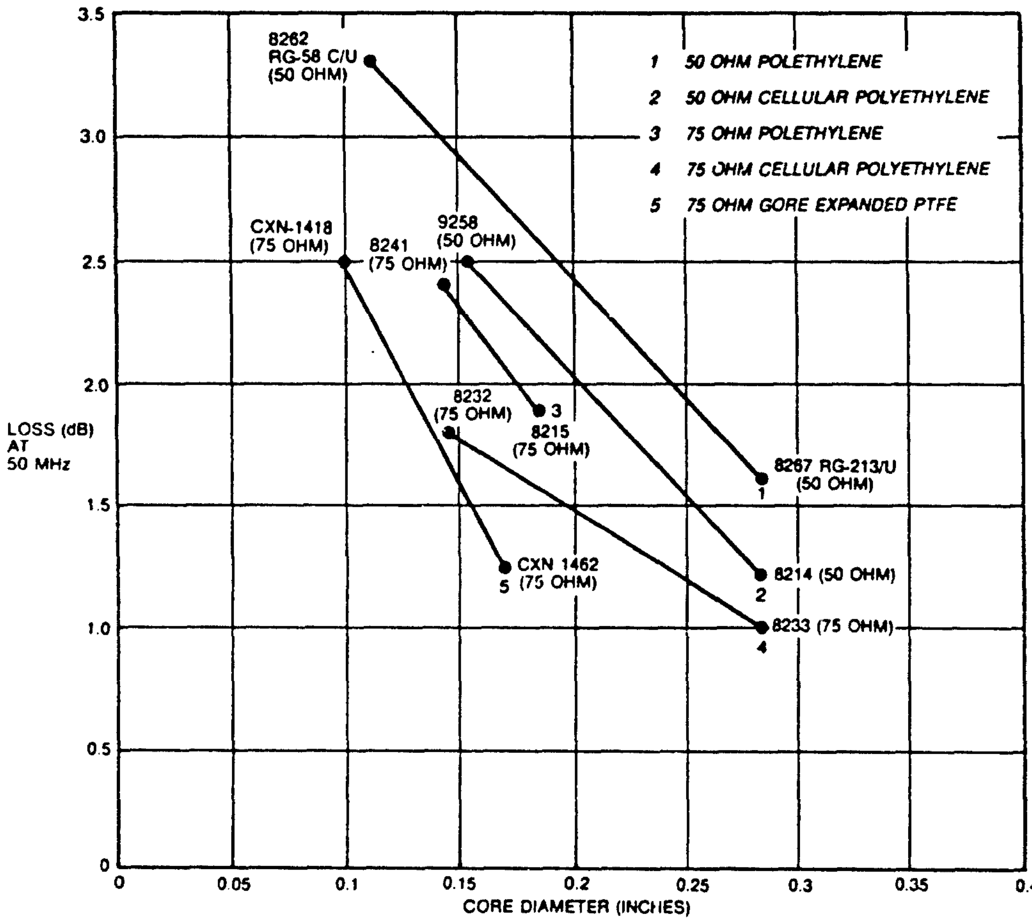

Figure 10. Impact of Size, Impedance and Dielectric on Attenuation

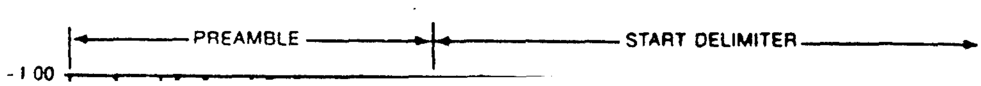

2. The dynamic range at the receiver was forced to accommodate the differential loss expected in a worst-case network which is greater by the amount of the network differential loss variation than would be the case for a more broadband transmission media.

Connectors

Because of the large number of nodes there may be as many as 128 connectors in series on the bus (two for each bus coupler). Two connector characteristics are, therefore, of importance: (1) impedance and (2) insertion loss. MIL-C-39012/1E specifies that the loss of an N-type connector is less than 0.15 dB at 10 GHz, and that at other frequencies the loss will be 0.05 $\sqrt{f(GHz)}$ dB. In the frequency range of interest the insertion loss is less than the 0.02 dB. Smaller connectors, such as TNC-type, will also meet this requirement over the frequency range of interest.

3.1.2 Determining Receiver Sensitivity Requirements

The receiver sensitivity requirement is driven by the need to maintain a S/N ratio adequate to allow the required error rate throughout the network. The worst case S/N ratio operating point occurs at the point of maximum separation between transmitter and receiver (59 dB) and while under worst-case EMI conditions. Figure 6 shows the probability of error in a digital system at various S/N ratios. At the design operating point for the HSDB, a S/N ratio of 19 dB is required (theoretical). Since theoretical performance cannot be achieved, an additional 6 dB was added to allow for transmitter and receiver induced errors. This results in a conservative system design point.

Theoretical S/N Ratio Required	19 dB
Design Margin	6 dB
Network Design S/N Ratio	25 dB

Noise at the receiver input port is attributable to several sources:

a. EMI radiation on the coax
b. Transmitter leakage
c. VSWR at component interfaces
d. Noise figure of the receiver

Each of these was investigated separately in order to arrive at an estimate of the level of noise expected at the receiver port. RG-196 coax was assumed for the Tx and Rx stubs; RG-142 coax was assumed for the mainbus. These were selected because they were to be used in the

demonstration equipment and because they were typical of cables found aboard aircraft. In installations using other types of cable, the analysis may need to be performed using modified variables. The result of this analysis is summarized in Figure 11. It shows that a minimum signal level of -35 dBm is required in order to maintain a 25 dB S/N ratio under worst case conditions. This appeared to be a practical design point so it became the receiver sensitivity specification for Task I.

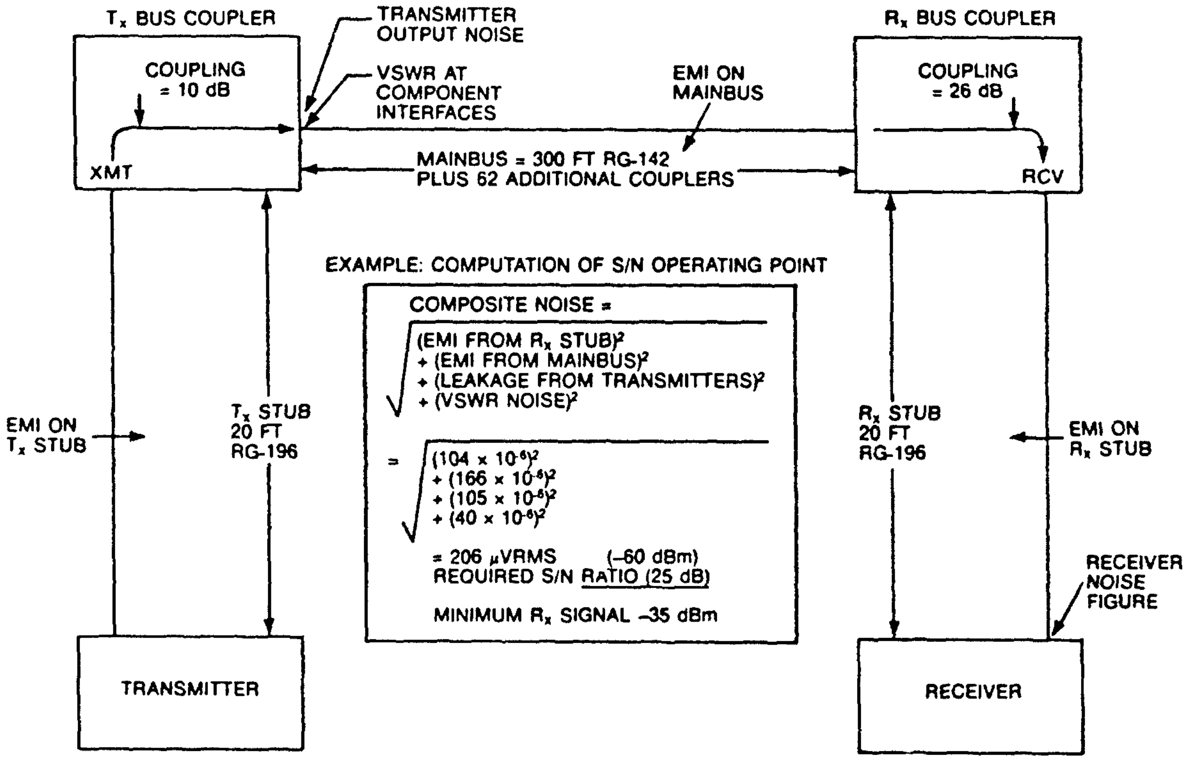

Figure 11. Summary of S/N Analysis for a Typical Coaxial Network

EMI Induced Noise

The broadband EMI noise we can expect to encounter was determined by analysis because no available data on the broadband noise in the frequency of interest could be found. The analysis was performed using Faraday's induction law. Figure 12 illustrates the conditions under which the analysis was performed.

Assuming that the main cable passed near 100 other electronic equipment, all radiating in phase the maximum field strengths allowed by MIL-STD-461B (0.1 Volts/Meter/MHz each, 10 Volts/Meter/MHz total and a uniform spectrum from 10 kHz to 100 MHz), the induced noise on the main bus would be 417 μV RMS. Assuming that each stub may be subjected to radiation from 25 pieces of equipment, the induced noise on the stub would be 166 μV RMS.

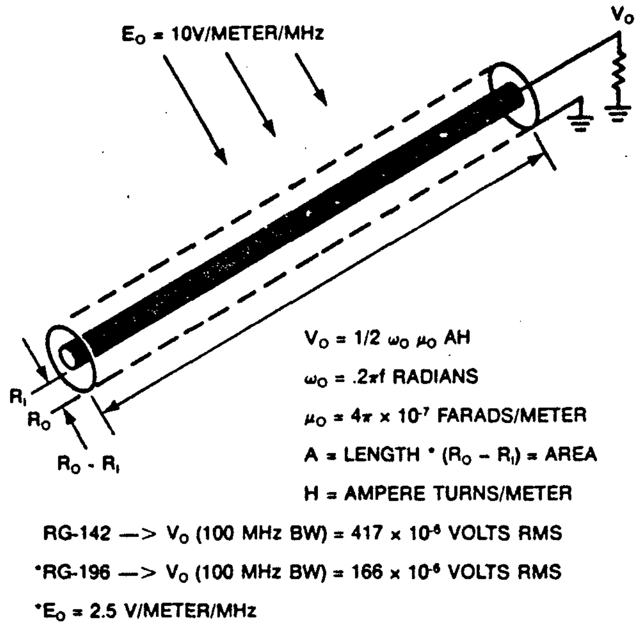

Figure 12. EMI Induced on a Length of Coaxial Cable

Transmitter Leakage

Transmitter leakage is specified as a maximum of 300 μV RMS on the mainbus, per transmitter. The transmitter noise contribution peaks at the center of the mainbus since noise contributed from both ends of the network is most additive at that point. The S/N ratio is worst at the end point of the network, however, because the signal level is much higher at the center of the network. This point was verified by a computerized simulation of various network configurations. At a receiver located at the end point of the 64-node network the composite transmitter leakage is 105 μV RMS. This became the worst-case design point for the network.

VSWR at Component Interfaces

Reflections caused by impedance discontinuities along the mainbus appear as intersymbol interference at the receiver. This form of interference was treated as an additional noise component during the analysis. The noise profile across the network was found to be similar to that of transmitter leakage. At the end point of a 64-node network, the composite VSWR noise is 40 μV RMS at the receiver input.

Noise Figure of the Receiver

Performance of the network is principally dependent upon the noise present at the receiver input. Due to the relatively high level of signal present, even under worst-case conditions, the receiver noise figure was found to be not significant and was ignored after the initial analysis.

3.1.3 Determining Transmitter Power Output Requirements

Determining the power required from the transmitter consisted simply of consolidating the receiver sensitivity operating point with worst case network loss as follows:

Receiver Sensitivity	-35 dBm
Network Loss (max)	59 dB
Required Tx Output Power	+24 dBm

The figure, +24 dBm, represents a substantial power, 250 mW, but is within the limits of practicality and was therefore selected as the specified transmitter output power for Task I.

3.1.4 Determining Receiver Dynamic Range Requirements

Determining the dynamic range requirement for the receiver consisted of defining what signal level would be seen by the receiver when it was the transmitting node compared with the receiver sensitivity operating point.

Tx Power Out	+24 dBm
Network Loss (Min)	36 dB
Max Rx Signal Level	-13 dBm
Min Rx Signal Level	-35 dBm
Rx Dynamic Range	23 dB

Since component variation will potentially cause individual receivers to see slightly greater signal ranges, a design margin of 3 dB was added to arrive at the specified value of 26 dB.

3.1.5 Network Operating Envelope

The HSDB network design arrived at by the procedure just described satisfies the requirements of Rockwell's contract. It represents a single point design. Obviously, other network configurations are practical, within some envelope. For example, operating with low loss coax will allow the network to operate with greater than 100 meter separation or with greater than 64 couplers. Achieving a coupler design with lower loss will provide a similar extension of the operating envelope. For example, with 0.1 dB couplers interconnected using Belden 8233 cable could satisfy most Air Force airborne applications up to 1600 feet in length. The Belden 8233 is a double shielded cable approximately 0.475 inches outside diameter. To further extend the operating envelope, a lower loss cable, such as Comm/Scope P3-75-500J, is required. With

Comm/Scope P3-75-500J, a system well over 1000 feet in length, with 128 terminals, and 0.1dB couplers, could be accommodated.

3.1.6 Modulation Characteristics

The modulation (encoding) method was selected to provide the required level of performance on the selected media (cable). Characteristics of interest include: 1) it should occupy minimum bandwidth (this minimizes the loss differential across the operating bandwidth), 2) it must provide for local clock recovery, and 3) together with the detection method, it should be tolerant to other media effects.

A number of alternative modulation methods were considered. Among them were Manchester bi-phase, MSK (Minimum Shift Keying), and Bipolar NRZ. Bipolar NRZ was quickly rejected because the clock cannot be recovered easily. Both Manchester bi-phase and MSK allow the clock to be recovered from the data. Manchester bi-phase, the same method as used for MIL-STD-1553B, has a transition in the signal in the middle of every bit time. MSK is a phase continuous frequency shift keying method. Neither offered any particular advantage so far as ease of clock recovery was concerned.

One of the serious problems with wideband baseband systems is the potential for excessive gain variation, from dc to nearly two times the data rate, encountered by the receiver at the far end of the cable. The spectrum associated with these two alternative approaches is shown in Figure 13. Note that the spectrum for Manchester modulation has significantly lower gain variation. For this reason, it was selected for the HSDB.

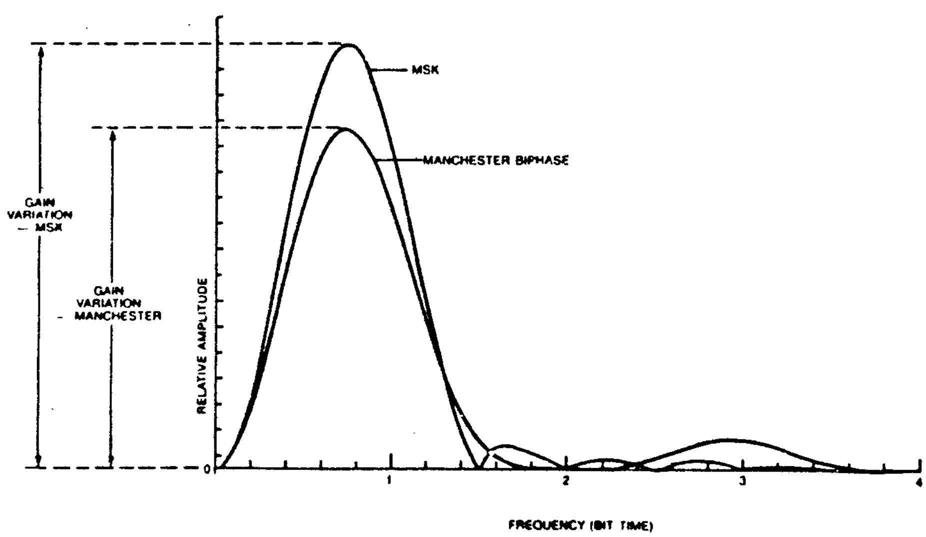

Figure 13. Spectrum of MSK and Manchester Biphase Modulated Signals

3.2 Coupler Development

As described earlier, engineering work during Task I focused on development of a superior coupler design. Achieving minimum excess loss and well controlled impedance and group delay characteristics were the primary design objectives. This was accomplished while maintaining a passive design which isolated stub and transmitter/receiver faults from the mainbus.

In operation, the bus coupler performs two functions:

1. It couples power applied to the coupler Tx port to the mainbus ports, driving both directions along the mainbus.
2. It couples power from the mainbus, arriving from either direction, to the coupler Rx port.

One significant design consideration was to minimize intersymbol interference caused by impedance discontinuities in the mainbus. The bus coupler was designed to function as a section of the mainbus with a characteristic impedance equal to the characteristic impedance of the mainbus (50 Ohms). It must have a low insertion loss, a low VSWR, and be fault tolerant. As a design goal, the coupler was to have less than 0.1 dB insertion loss, the VSWR was to be less than 1.035:1. Also, any failure in the bus coupler or transmit/receive unit could not cause the rest of the bus to be inoperable.

In many installations it is desirable to have long Tx and Rx stubs between the coupler and the terminal. This allows large physical separation of the redundant buses. To have long stubs, the receiver impedance must be matched to the stub. In order to minimize the impact of shorts on the stubs the coupler must step the reflected impedance on the bus up to some high value. For our design an impedance of greater than 3500 Ohms was chosen. A passive receiver coupler that satisfies those requirements is shown in Figure 14. The SPICE equivalent circuit is shown in Figure 15. The two transformers in series appear as a resistive load shunted by a capacitor (stray capacitance).

In the transmit mode, the coupler must match the stub characteristic impedance to one half of the characteristic impedance of the bus. (This matches the mainbus load, Z_O to the left in parallel with Z_O to the right.) Also the transmit stub should not be connected to the mainbus during the receive mode. This prevents noise generated by the transmitter from appearing on the mainbus. A transmit bus coupler that can decouple the transmitter from the bus except when transmitting is shown in Figure 16. The SPICE equivalent circuit of the transmitter is shown in the transmit mode in Figure 17, and in the receive mode in Figure 18.

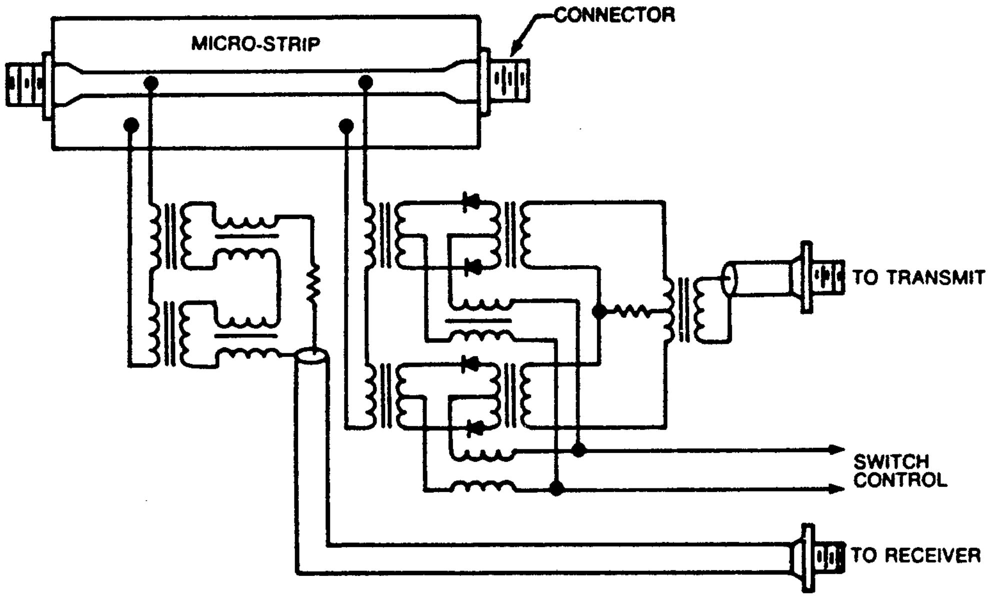

Figure 14. Simplified Schematic Diagram of the Coaxial Network Coupler

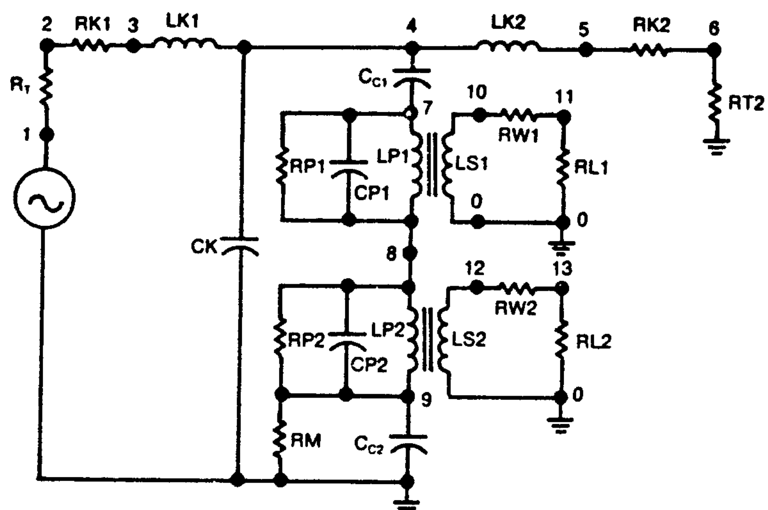

Figure 15. Receive Port Coupler Equivalent Circuit

Figure 19 models the coupler, both transmitter and receiver functions, from the perspective of a section of the mainbus. The accumulated capacitance, CP1 and CP2 in Figure 15, and CP1 and CP3 in Figures 17 and 18 combine to form the shunt CK shown in Figure 19. The shunt impedances RP1 and RP2 of Figure 15 and RP1 and RP3 of Figures 17 and 18 combine to form "G" of Figure 19. Careful mechanical and electrical design were used to make the coupler appear to be a good approximation of a 50 Ohm transmission line. Designing good

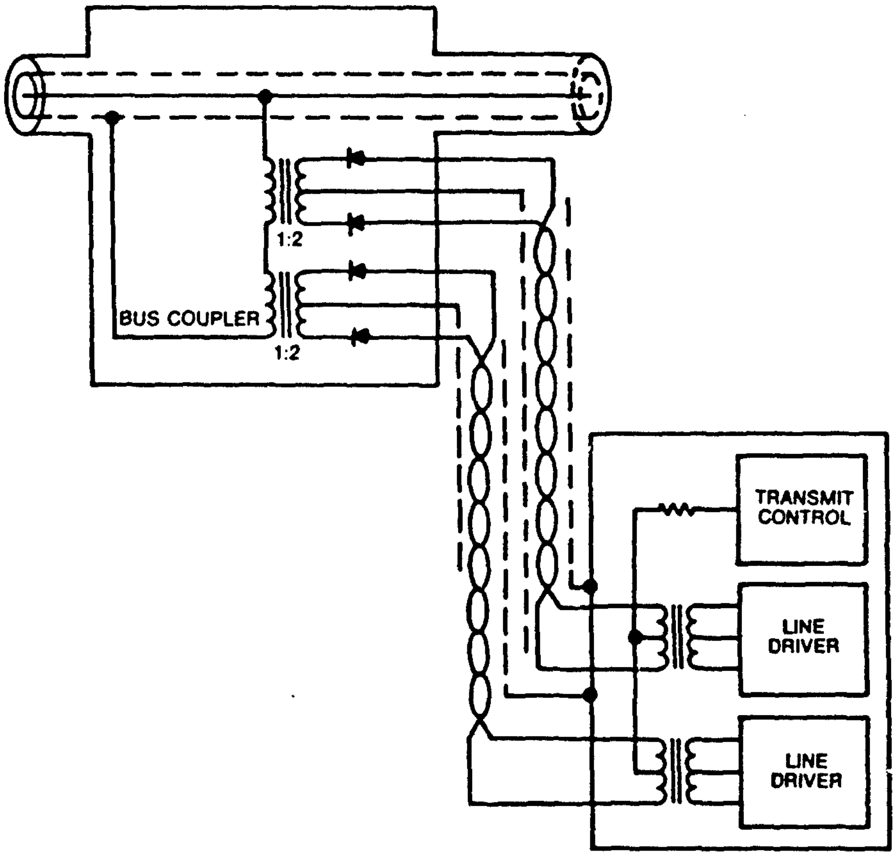

Figure 16. A Dual Diode Switch Disconnects this Transmitter Port from the Main Bus

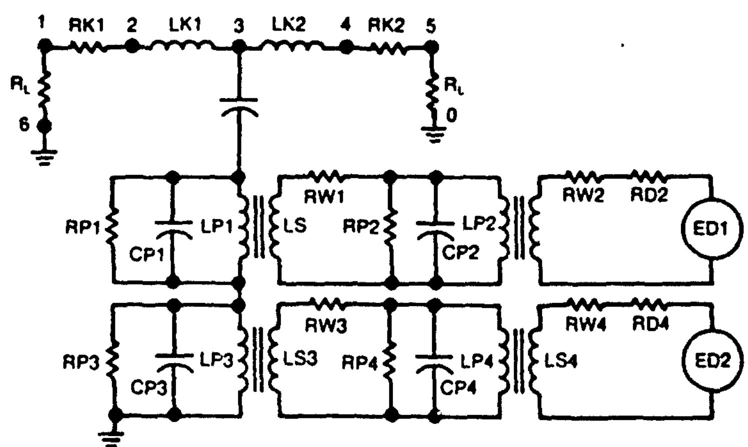

Figure 17. Transmitter Port Equivalent Circuit in Transmit Mode

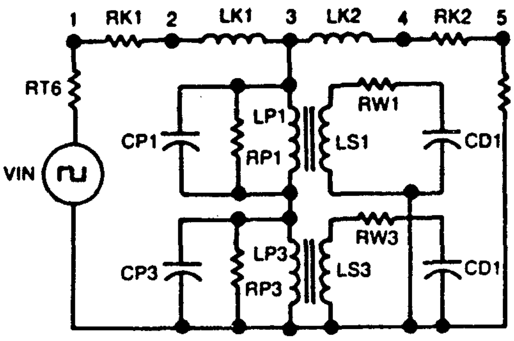

Figure 18. Transmitter Port Equivalent Circuit in Receive Mode

transformers and maintaining a low VSWR (1.1:1) were the most difficult design problems. The design goal was achieved, however, and verified by the fabrication of 64 couplers on a pilot production line. Reproducibility proved to be excellent with all couplers meeting specification after only nominal alignment. Figure 20 shows one of the prototype couplers.

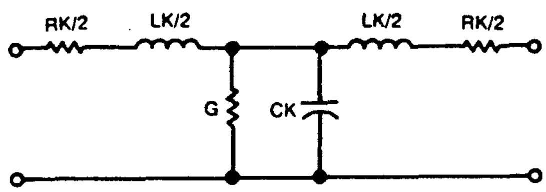

Figure 19. Lumped Constant Transmission Line Section

3.3 Transmitter Design

Design of the transmitter proved to be a relatively straightforward engineering task. The requirement was to accept an ECL level pulse train which comprised the Manchester encoded data message and amplify/condition it to produce a +23 dBm equivalent signal into the 50 Ohm Tx stub. After reviewing several alternative approaches, the design shown in Figure 21 was selected. A commercially available hybrid power amplifier produces an output directly compatible with the network. It is driven from a pair of ECL gates, through a transformer. The gates provide sufficient drive for the power amplifier and also provide an enabling logic input function.

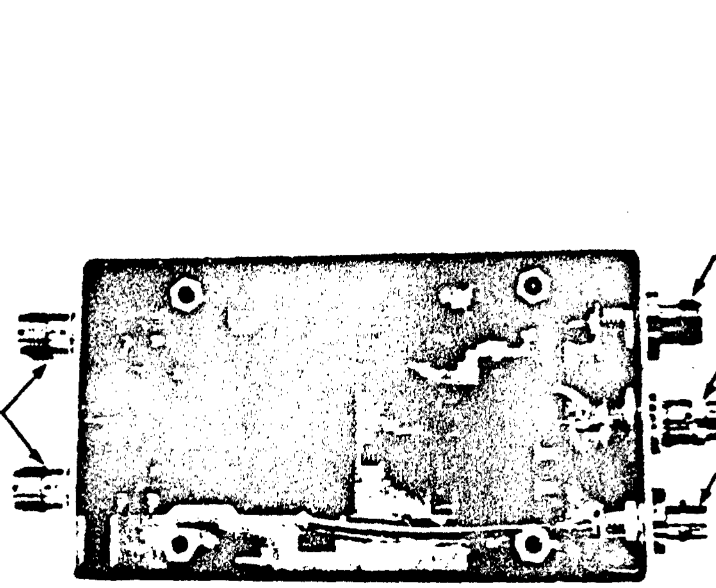

Figure 20. Prototype Wire Bus Coupler

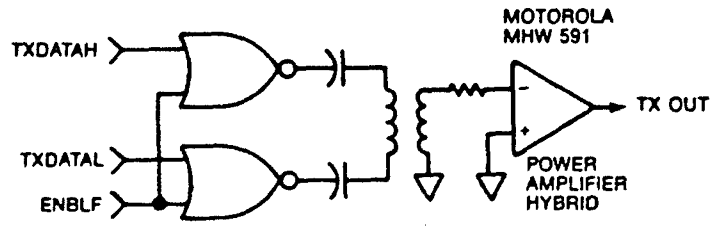

Figure 21. The Coaxial HSDB Transmitter uses a Hybrid Circuit Power Amplifier

3.4 Receiver Design

Design of the receiver proved to be somewhat more complex than anticipated, primarily because of distortion on the received signal which was attributable to frequency rolloff characteristic of the network. The primary source of distortion on the coaxial cable is the amplitude response as a function of frequency. This is shown graphically in Figures 22 and 23. Figure 22 shows a 4-bit pattern (1001) of Manchester encoded square wave signal as might be generated by the transmitter. Figure 23 shows a SPICE simulation of how that same pattern would appear at the end of the network. Note the severe distortion of the high frequency component of the waveform. The peak-to-peak amplitude of the first zero is 50 mV. This is less than the amplitude of overshoot which can be expected at the receiver port of the transmitting coupler. From this it is apparent that the method used by the receiver to detect the signal is very important if reliable reception is to be realized. A simple amplitude threshold detector such as is commonly used for MIL-STD-1553B, is not adequate. Instead, the method of detection must be relatively insensitive to amplitude distortion.

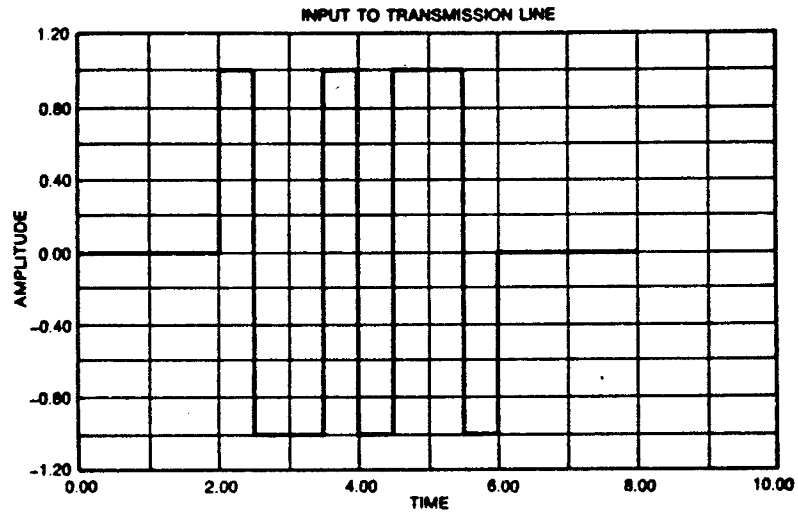

Figure 22. Manchester Waveform for '1001' Bit Pattern

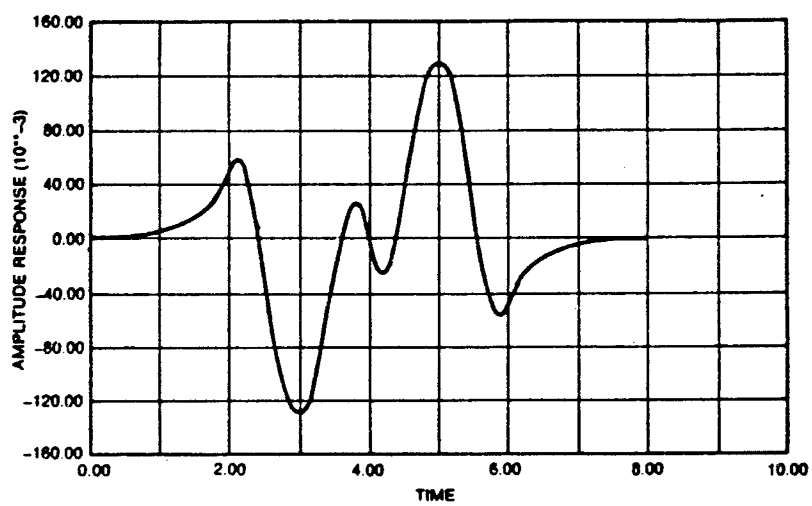

Figure 23. Simulated Degration of a '1001' Bit Pattern
at the End of the Network

Several different types of amplitude insensitive data detection methods were considered the most promising were: (1) zero crossing with time state, (2) simple phase detection, and (3) integrate and dump. A method equivalent to the integrate and dump type appeared to hold the best promise. The integrate and dump time interval is for a one half bit time period. In this scheme, the clock must be synchronized to the received data stream.

Figure 24 shows the receiver functional block diagram. The signal from the network is first amplified and then directed to both the clock extractor and the data recovery circuit.

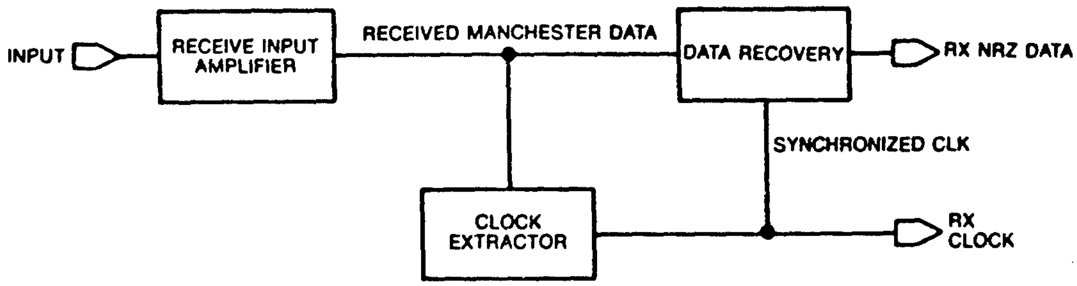

Figure 24. Receiver Functional Block Diagram

Clock Extractor

The purpose of the clock extractor is to recreate the 2X (100 MHz) clock which was used at the transmitter to encode the NRZ data into Manchester format. The objectives are (1) to produce (acquire) a local clock signal which is frequency and phase synchronous with the transmitted clock, (2) create the local clock in a minimum time period, and (3) maintain a stable local clock after acquisition. These are contradictory objectives from an engineering perspective so a tradeoff was required to arrive at a design which was adequate. Two methods of performing clock extraction were investigated. The first, a direct phase selection approach offered many advantages in speed of acquisition and stable tracking but proved too layout sensitive to be practical when fabricated using discrete components. The second approach, and the one selected for the breadboard, was the ringing tank approach.

Figure 25 shows the ringing tank simulation circuit used for the analysis. The active stages represents any of several monolithic RF amplifier chips available commercially. The amplifier is used to drive an L-C tank circuit (L1, C9). Simulation (see Figure 26) showed that a usable clock could be achieved in four bit-times. The design was validated using breadboard hardware; the results are shown in Figure 27.

Data Recovery

As described earlier, several different methods of data recovery were investigated in order to arrive at a design which worked well with the distorted waveforms present on the coaxial HSDB. A very good design was arrived at after much SPICE simulation and breadboard testing. It used a complex detection algorithm which included:

a. 80% dependent on wavefront detection
b. 20% dependent on amplitude detection
c. Threshold hysteresis

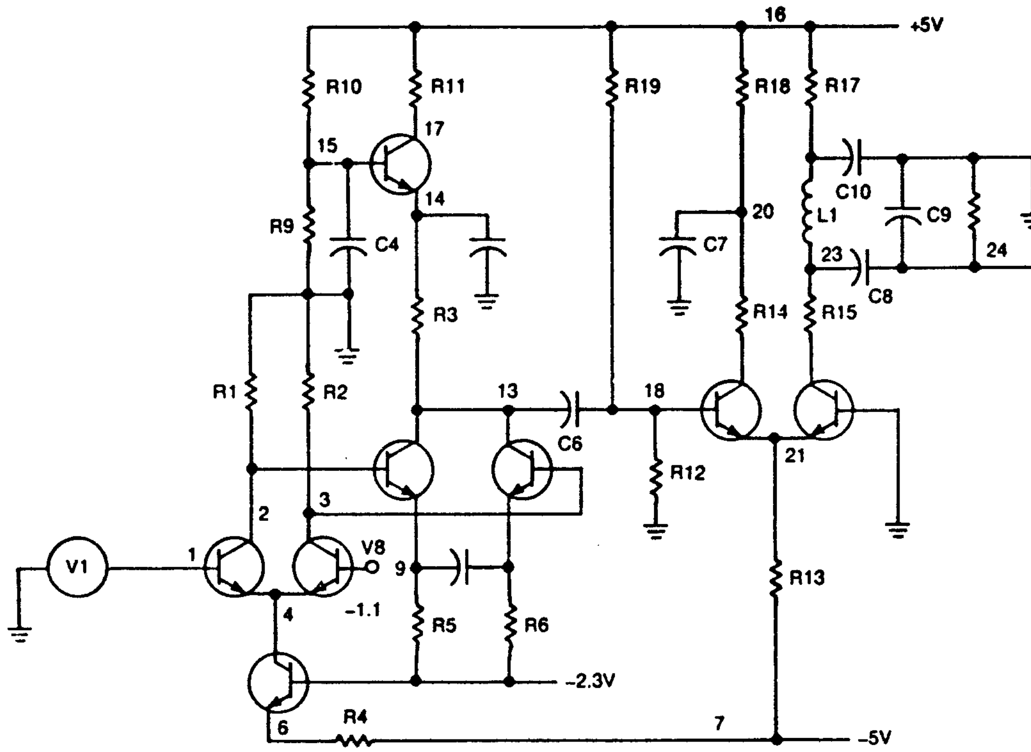

Figure 25. Circuit Used for Ringing Tank Simulation

This allowed signals from high amplitude (near) transmitters with significant waveform overshoot components and signals from low amplitude (far) transmitters with severe high frequency rolloff distortion to be received reliably. Figure 28 shows the detector circuit used. It consists of a high speed comparator with both inputs fed from the single ended output from the preamplifier. One of the inputs is delayed 5nS from the other creating a condition wherein a phase reversal lasting 5nS is detected as a change in logic level.

Figure 29 is a SPICE simulation showing how the detector operates on high amplitude input waveforms having significant overshoot. Figure 30 is a similar illustration of how the detector operates on low amplitude signals having significant rolloff of the high frequency components.

3.5 Proof of Concept Testing

Designs for the coaxial HSDB network were validated using various breadboard and brassboard test configurations. All critical circuits were breadboarded and tested over temperature prior to being integrated into the breadboard transmitter/receiver units (TRUs).

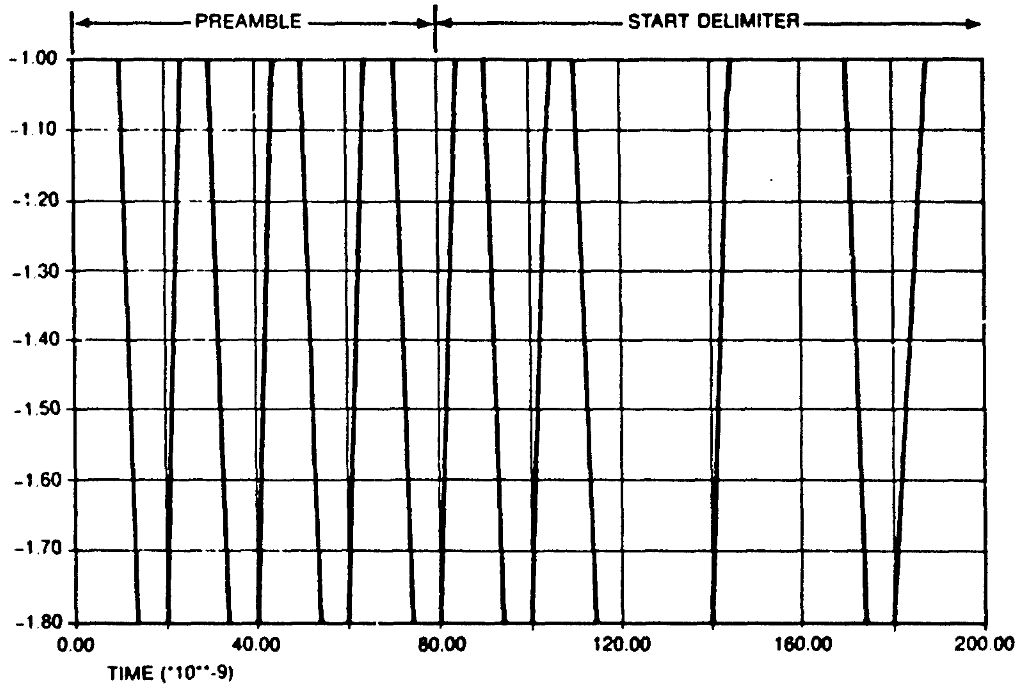

Figure 26a. Simulation Input Waveform

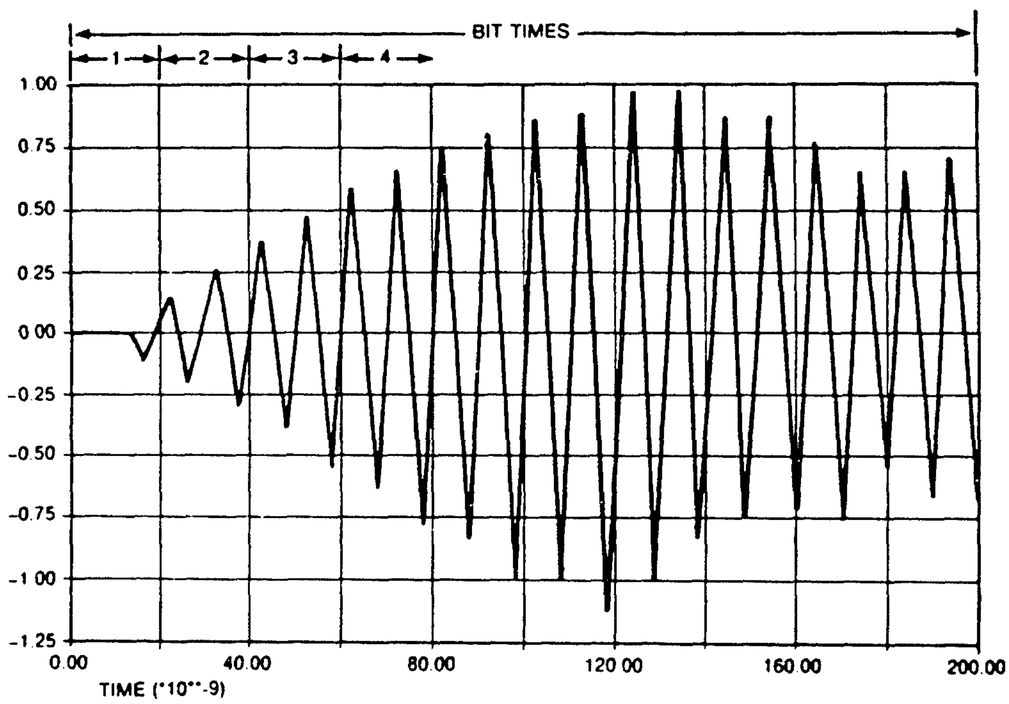

Figure 26b. Simulation Output Waveform

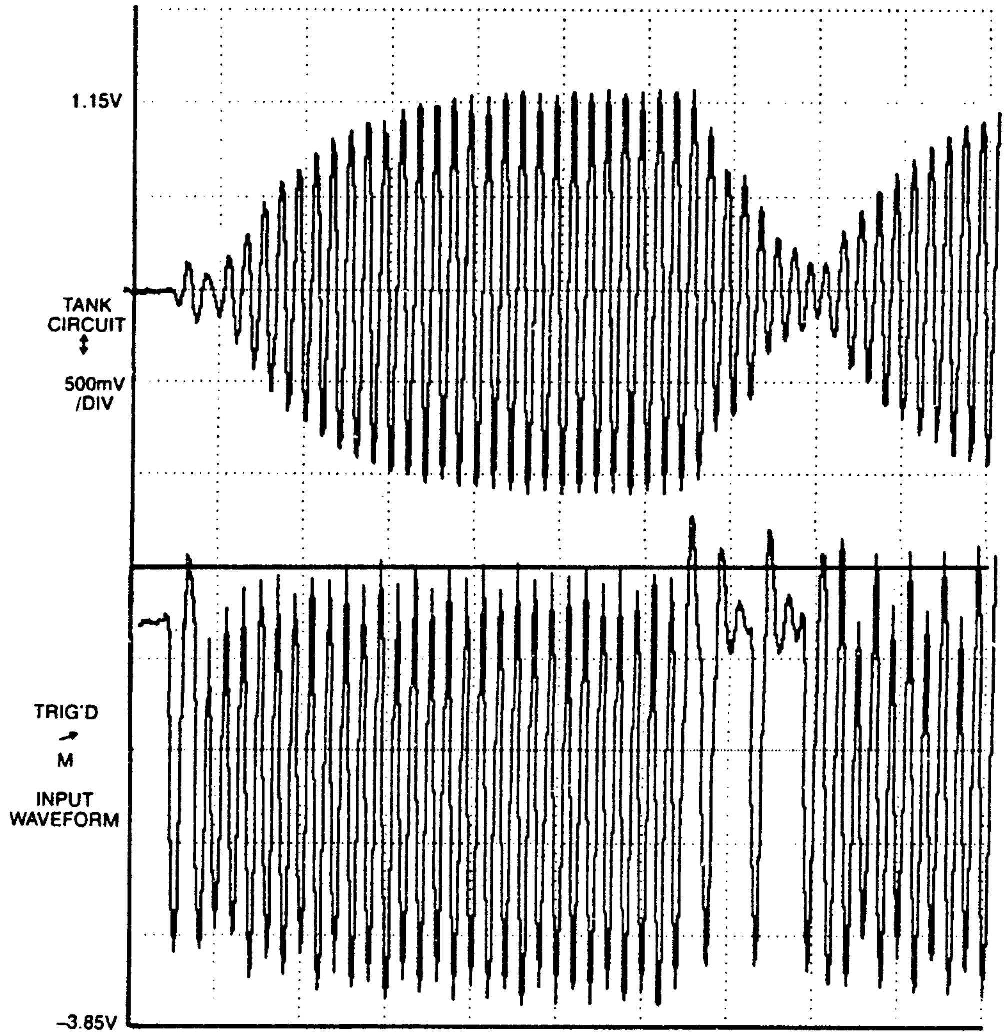

Figure 27. Breadboard Test Waveforms Validated Ringing Tank Clock Recovery Approach

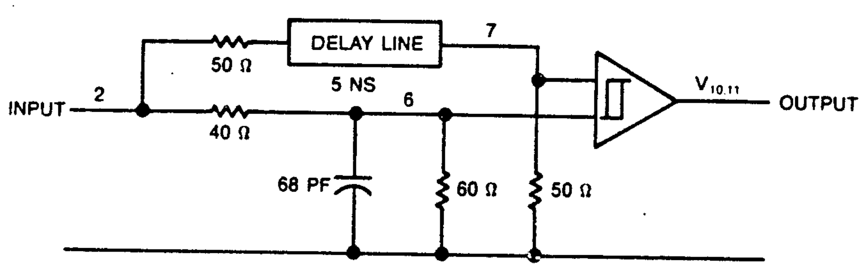

Figure 28. Detection Circuit

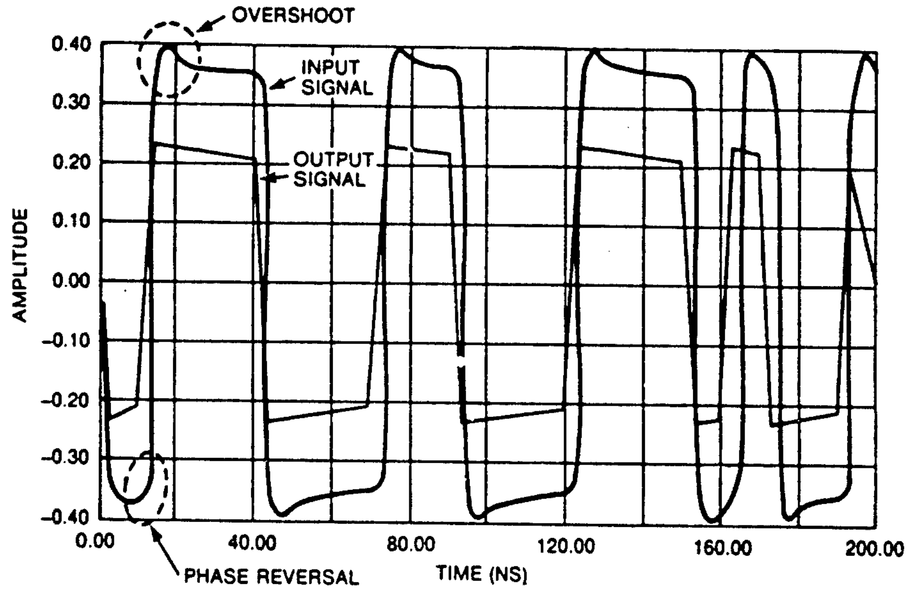

Figure 29. Detector Performance on 'Near' Signals

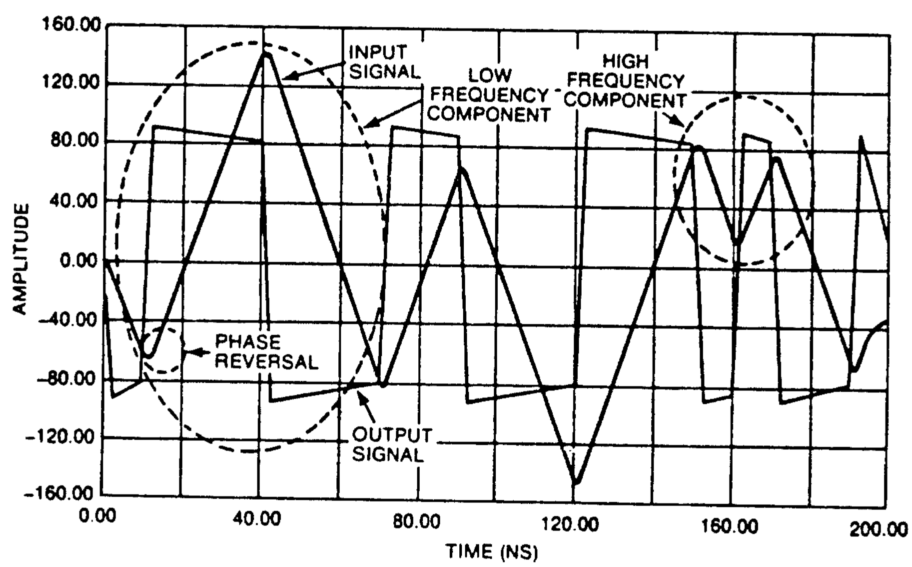

Figure 30. Detector Performance on 'Far' Signals

Each breadboard TRU consisted of two assemblies, (1) a transmitter circuit board and (2) a receiver circuit board. Two of each were built and tested in a simulated network configuration. Finally, six each brassboard transmitter and receiver circuit boards, and 64 brassboard couplers were fabricated, tested, and characterized. The test and characterization equipment designed for

this purpose is described in Section 6 of this report. Figure 31 is a photograph of the brassboard receiver circuit card; Figure 32 is a photograph of the brassboard transmitter circuit card.

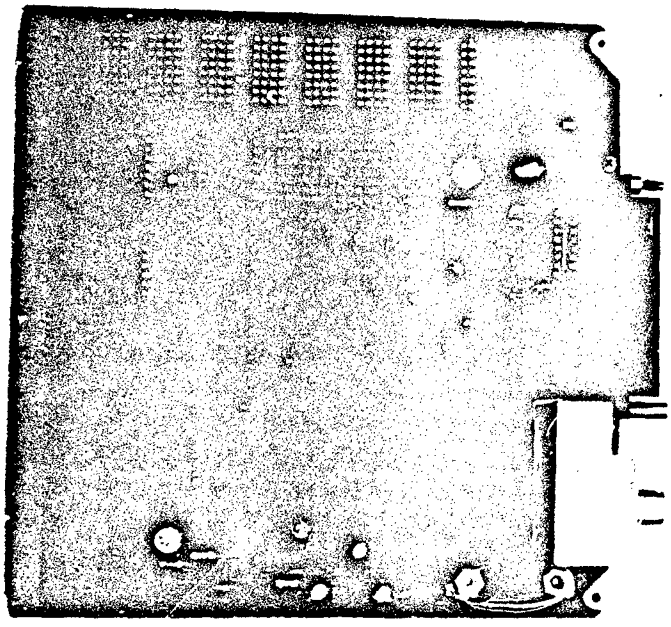

Figure 31. Brassboard Coaxial HSDB Receiver Card

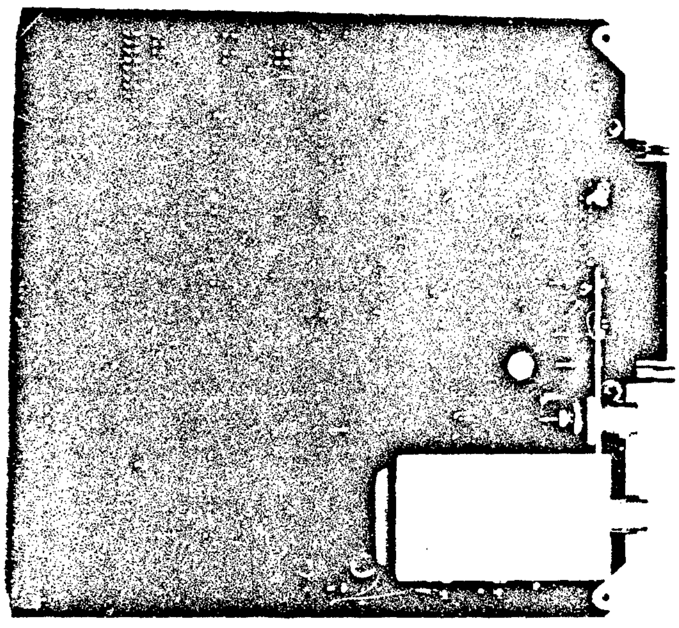

Figure 32. Brassboard Coaxial HSDB Transmitter Card

Characterization

Characterization is defined as the process by which parametric performance of the design is measured. The following characterization testing was performed on the six brassboard transmitters and receivers:

a. Transmitter output power
b. Transmitter output waveform
c. Transmitter clock stability
d. Transmitter synchronization waveform
e. Transmitter timeout override
f. Transmitter switch waveform
g. Transmitter output noise
h. Transmitter modulation
i. Receiver acquisition range
j. Preamble response time
k. Receiver dynamic range
m. Receiver input impedance
n. Bit Error Rate (BER)

Performance of the TRU was proven over the temperature range of -54 ^{o}C to +95 ^{o}C. The BER characteristics of a typical TRU is shown by Figure 33.

Demonstration

A demonstration of coaxial network HSDB technology was conducted using three brassboard TRU connected into the HSDB system demonstration equipment. This demonstration showed three HSDB terminals operating in a simulated HSDB network. The network was emulated using 64 brassboard couplers interconnected using random lengths of RG-142 coax with a total length of 100 meters. Figure 34 shows the HSDB demonstration equipment as configured for the Task I ATR demonstration. Figure 35 shows a typical waveform monitored on the network. Signals from the three different terminals are identifiable as three distinctly different amplitudes in the photograph. This occurs because of the differing length (and attenuation) of the network between the monitor and each transmitter.

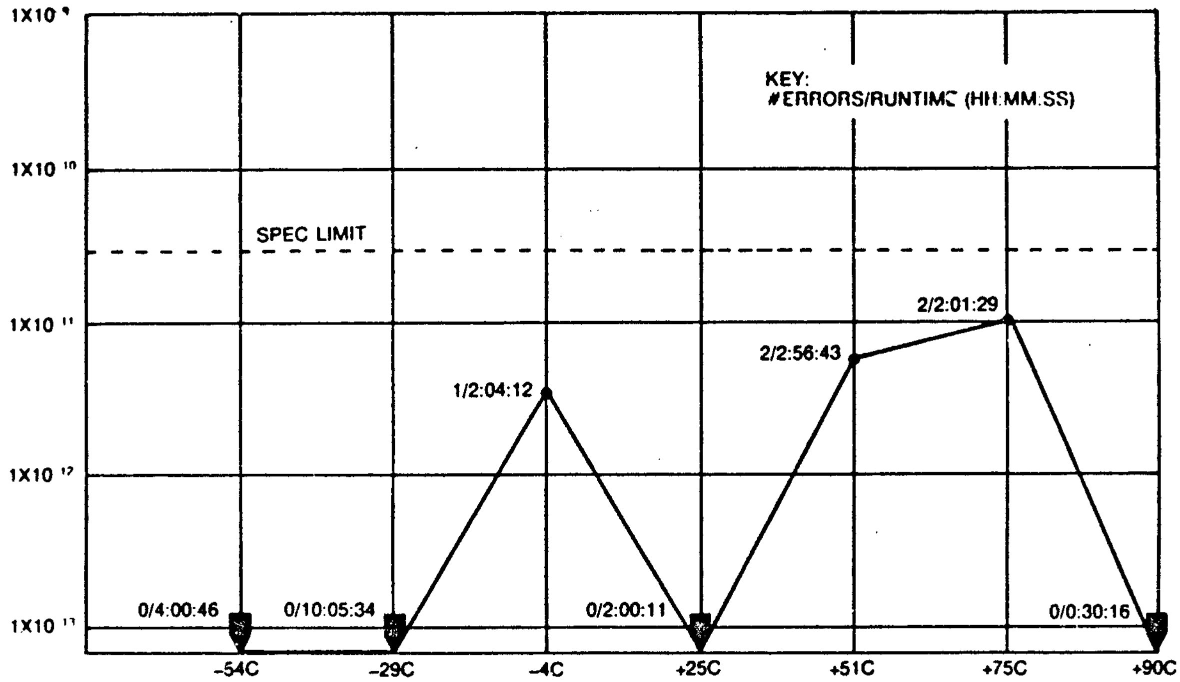

Figure 33. Brassboard Coaxial T/R Unit BER Characterization

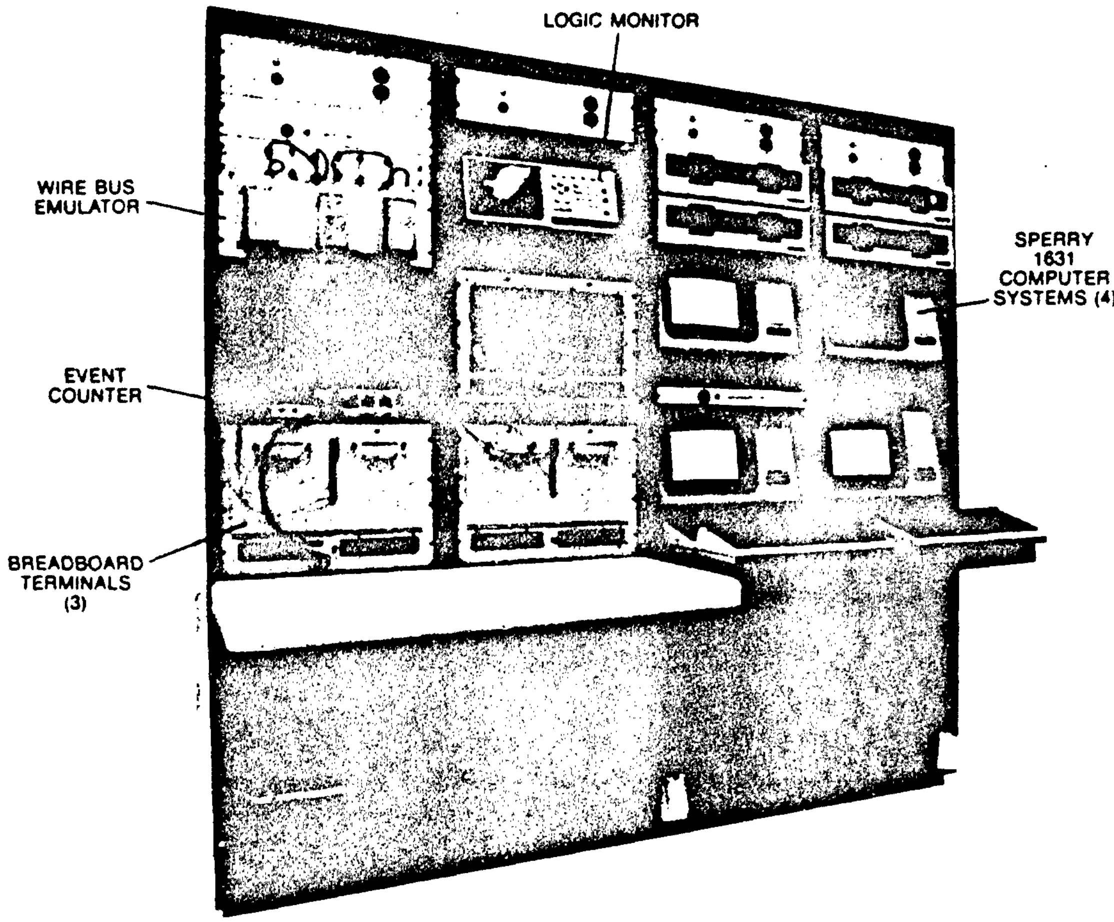

Figure 34. HSDB System Demonstration Equipment Configured with Coaxial T/R Units and Couplers

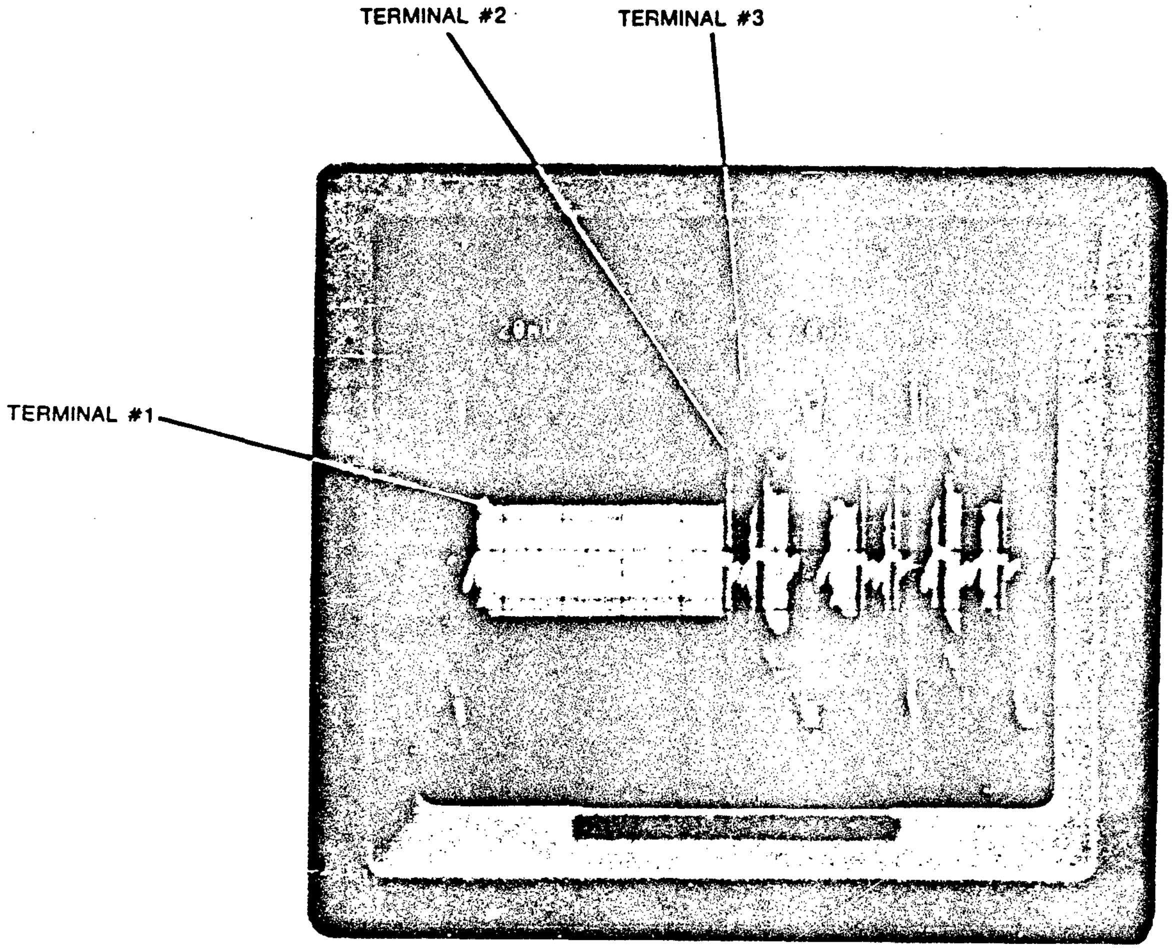

Figure 35. Typical Signals Monitored on the Coaxial HSDB System Demonstration Equipment

4.0 DEVELOPMENT OF FIBER OPTIC NETWORK TECHNOLOGIES

Development of the fiber optic version of the HSDB encompassed investigation of the technological limitations of fiber optics, establishing performance goals for the network, and design/fabrication/test of breadboard and brassboard hardware as proof of concept. This was performed as Task II of the HSB Technology Development Program as defined in paragraph 4.2 of the SOW. The majority of this effort was performed by FiberCom, Inc. of Roanoke, Virginia under subcontract to Rockwell.

Two implementation approaches to a broadcast topology network are possible. As shown in Figure 2, these are the linear bus and the star. The linear bus approach is preferred from the standpoint of installation and growth flexibility. The limited power budget allowed by present generation optical source and detector technology renders this approach impractical for the HSDB application, however. Instead, the single passive star topology was selected for design, development, test, and demonstration of the fiber optic HSDB. This does not imply that Rockwell recommends tne passive star topology for production aircraft installations; only that the objectives of this program could be met without the cost and risk associated with other alternatives which were more production oriented.

Finding a design which worked within the optical power budget allowed by the state-of-the-art in 1984 was our initial design challenge. The power budget is defined by several characteristics; (1) transmitter power output, (2) receiver sensitivity and operating range, (3) coupling loss, (4) coupler excess loss, (5) connector loss, and (6) fiber loss.

Transmitter power output and receiver operating range are limited by technology. Coupling loss is set by the number of nodes in the network since transmitter power must be divided among them. Coupler excess loss is determined by the coupler manufacturing process. It represents power which enters the coupler but does not exit on one of the output fibers. Connector loss and fiber loss are installation sensitive. Short runs with few connectors may exhibit little or no loss, long runs with many connectors may have in excess of 10 dB loss. It became apparent early during the program that achieving a superior receiver design was key to success of Task II. Engineering work focused heavily in this area. The result was development of a fiber optic receiver with state-of-the-art sensitivity and dynamic range performance. It also became apparent that a purely passive network could support a very minimal interconnect. This, in effect, dictated that the interconnect used for the demonstration would not be representative of production aircraft installations.

LED sources at 850 nm wavelength were selected over laser diodes because they are much easier to drive. Multimode fiber was selected for the interconnect because of its ability to easily couple power from LEDs and through connectors. Networks using step-index fiber are a practical alternative, dependent upon the installation characteristics of any specific application.

4.1 Network System Design

As shown in Figure 36, the fiber optic HSDB network consists of transmitters, receivers, a coupler, and fiber interconnect. The system optical power budget is a principal technological driver in a network of this topology. The design for the fiber optic HSDB network evolved through the two-step process described below:

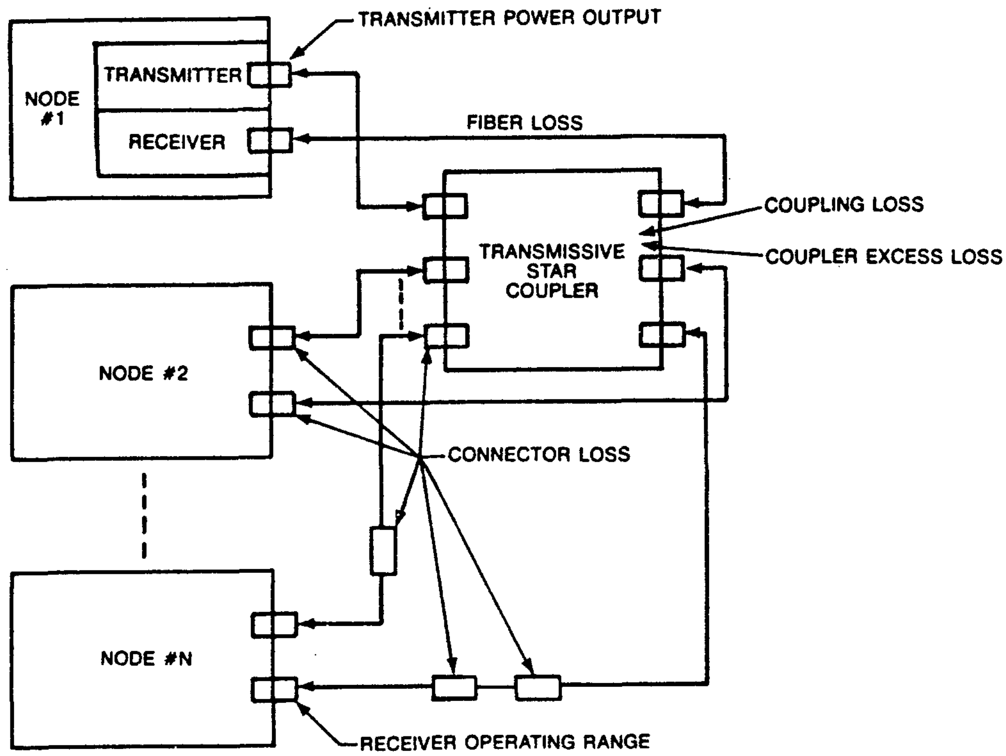

Figure 36. The Fiber Optic HSDB Topology is a Single Transmissive Star

a. First, an analysis of available technologies was performed in order to allocate system characteristics, most notably optical loss, among the elements comprising the system. The analysis included topologies, couplers, fibers, connectors, sources, and detectors. From the analysis came the system level requirements which drove the Task II development effort. The analysis phase is described, in detail, in paragraph 4.1.1.

b. Next, a synthesis phase was initiated. The synthesis phase resulted in definition of specific engineering approaches and technologies which would be followed during development. Critical requirements for each element of the network were

determined and documented in the development specification. The synthesis phase is described, in detail, in paragraph 4.1.2.

Table 5 summarizes the requirements which resulted from the system design effort. These were treated as design goals for the ensuing development phase. In some cases, requirements were changed during the course of the prog: im as a better understanding of the capabilities and limitations of the technology was achieved. Final values are found in the HSDB system specification.

Table 5. Summary of Fiber Optic Network Specification Requirements

Receiver Operating Range: The receiver shall meet all performance requirements while receiving signals of greater than -35 dBm peak and less than -12 dBm peak.
Receiver Dynamic Range: The receiver shall meet all performance requirements while receiving transmissions wherein packets are preceded by 8 bits of preamble and 200 nS off time, and alternate packets are less than 21 dB different in amplitude.
Transmitter Power Output: Transmitter output power shall be greater than -6 dBm peak and less than 0 dBm peak.
Network Path Loss: Path loss from transmitter to receiver shall be less than 29.5 dB and greater than 10 dB at 850 nm.
Modulation Format: Modulation format shall be on-off Manchester at 50 Mbps.
Fiber: Fiber interconnect shall consist of 100/140 μm graded index multimode fiber with nominal numerical aperture of 0.29.
Wavelength: Wavelength shall be 850 nm, nominal.
Connectors: SMA style 906 or equivalent, less than 2 dB loss per connection.
Optical Source: LED at 850 nm wavelength
Optical Detector: Optical PIN diode
Topology: Single passive transmissive star

4.1.1 Analysis Phase

The goal of the analysis phase was to incorporate SAE AE-9B/L high speed bus standard and PAVE PILLAR requirements into the HSDB system. This approach was selected to encourage development of a HSDB which was applications oriented rather than merely a technology demonstration project. Two principal topics were addressed in the analysis, optical power budget and topologies. A review of the state-of-the-art of the various elements of a fiber optic data bus was conducted. Based on the characteristics of the various components, an analysis of data bus system performance was performed. The bus elements considered and the factors evaluated are shown in Table 6. In addition to the factors listed, cost was considered for each element.

Table 7 summarizes the results of the analysis phase. It should be noted that all elements of the analysis were accomplished in an interactive and approximately parallel process. The discussion which follows may imply a serial process, where one element was studied and a requirement established prior to beginning the next phase. Such was not the case and knowledge of the actual process may add clarity to some of the discussions which follows.

Table 6. Technology Drivers

COMPONENT	FACTORS
Couplers	Losses Number of Taps
Cabling	Fiber Type Connector Splicing
Optical Source	Power Speed
Optical Receiver	Sensitivity Operating Range Intermessage Dynamic Range Clock Recovery
Processing/Interface Logic	Speed Power Consumption
Topologies	Performance Reliability Flexibility

4.1.1.1 Power Budget Analysis

A key element in the design and optimization of any fiber optic link, including the PAVE PILLAR HSDB, is the system power budget analysis. Such an analysis is important not only to ensure that there is adequate optical power at any given receiver under all conditions, but to also ensure that there is not too much optical power at any given receiver.

Three basic characteristics must be considered: (1) optical source output power, (2) optical receiver sensitivity, and (3) system losses. Each of these were studied independently in order to ascertain the state-of-the-art. Since the system losses are basically the same for either an LED or

Table 7. Fiber Optic Analysis Summary

	MIN	MAX
System Loss Calculation:		
Nominal loss for coupler (64 nodes)		18 dB
Nominal loss for coupler (8 nodes)	9 dB	
Coupler excess loss	0 dB	3 dB
Fiber interconnect loss	0 dB	2 dB
Connector Loss (1.5 dB ea)	0 dB	6 dB
	9 dB	29 dB
Transmitter Power Output:	-6 dBm	-3 dBm
Receiver Operating Range (ROR) Calculation:		
Transmitter power	-6 dBm	-3 dBm
System loss	29 dB	9 dB
Receiver Input	-35 dBm	-12 dBm
ROR		-23 dB

a laser source, the maximum allowable system loss (systems loss budget) was defined to be the difference between practical transmitter output and the sensitivity of a practical receiver design. Figure 37 illustrates the analysis results, and shows the established design goals for Task II development. Note that both receiver sensitivity and optical source power output are data rate sensitive. At 50 Mbps a LED/PIN diode transmitter/receiver design can accommodate about 30 dB system loss. The theoretical loss of a 64-node network is 18 dB. This allowed a design margin of 12 dB for interconnect, coupler excess loss, and engineering margin. This was thought to be sufficient to allow development of demonstration hardware at acceptable cost and risk. For this reason, our design points were tentatively set at -35 dBm for receiver sensitivity and -6 dBm for transmitter power output.

Optical Sources

Two alternatives were studied to select the HSDB optical source; these were the LED and the laser diode. Each has its advantages and drawbacks. In keeping with the applications oriented objective of the analysis phase, the following device characteristics were studied:

a. Output power
b. Ease of use
c. Bandwidth
d. Reliability
e. Cost

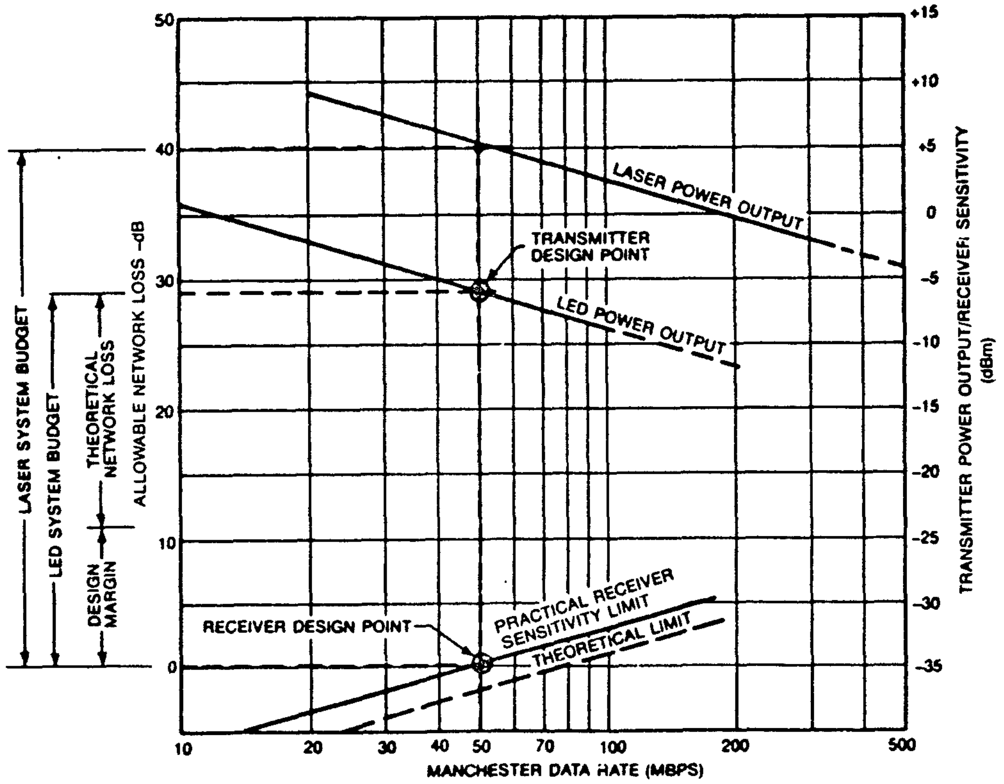

Figure 37. Summary of Fiber Optic System Constraints

The study was done assuming that only small improvements in performance could be expected during the initial application period. In effect, parts available at the time of the study were judged to be representative of performance attainable in early generation production systems.

Output power is clearly of importance since it relates directly to the allowable operating envelope of the network. Here the laser diode has a substantial advantage. This is illustrated in Figure 38. Although both devices are similar in construction, the laser diode is designed to operate in regenerative mode. This is shown in Figure 39. Below the knee a laser diode acts just as a LED, increasing the current through the junction will increase the light produced in a fairly linear fashion. When the diode reaches a point where laser action starts the light output vs current function is much more sensitive. This gives rise to many applications complexities.

One of the complexities of laser diode operation is that of temperature compensation. Note that the LED requires relatively small compensation of drive current to provide a constant output power over the entire temperature range. It is also quite non-critical meaning that inaccuracies of the compensation circuit will not result in damage to the LED, the only affect is a small variation in the optical power output. Also, variation of this characteristic among similar diodes is relatively small. This means that a drive circuit can be designed which matches the

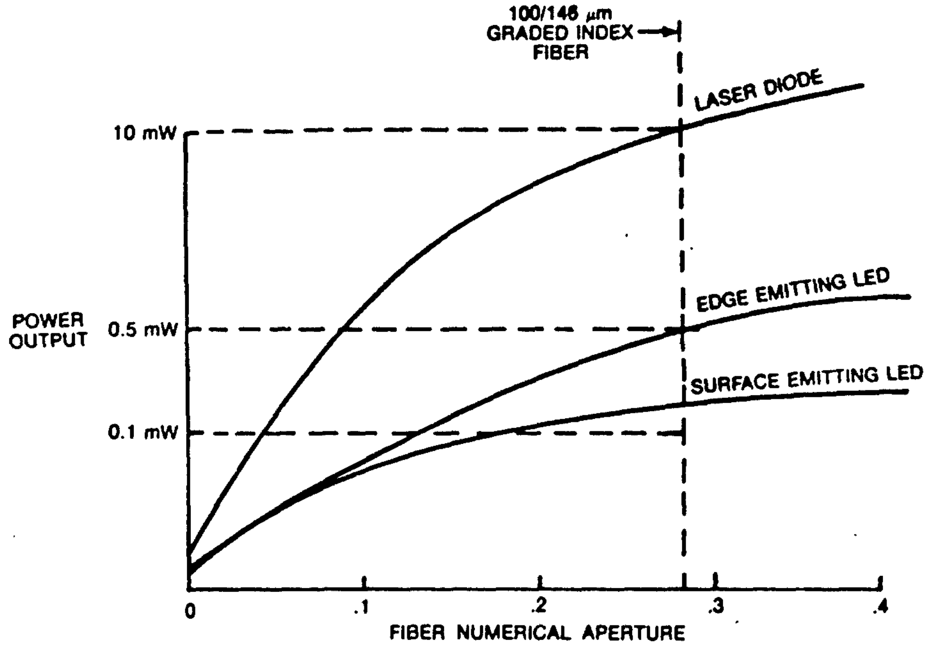

Figure 38. Power Output Capabilities of Different Sources at 100 MHz

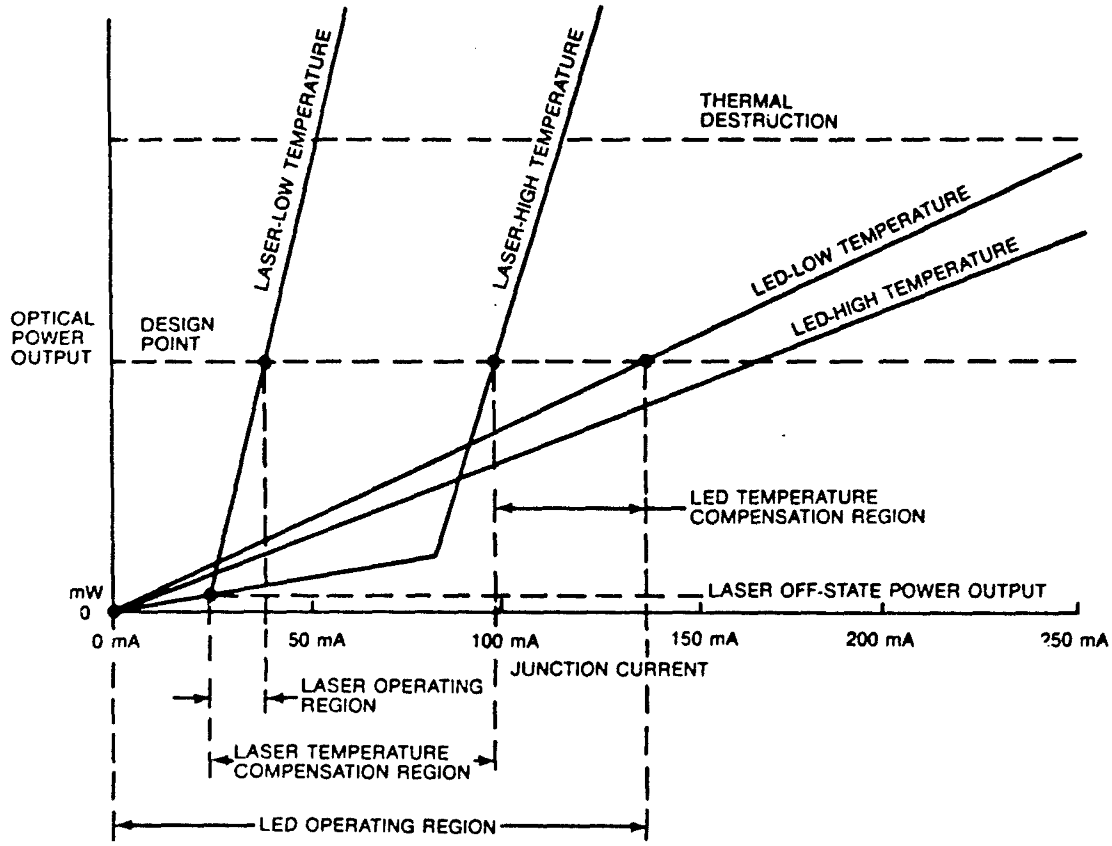

Figure 39. LED and Laser Output vs. Drive Current

characteristics of a group of similar LEDs. The laser diode is not so well behaved. Because of its regenerative nature, it is very sensitive to differences in drive current. In fact, it is relatively easy to drive it to destruction because of the steep slope of the curve. This characteristic is highly temperature sensitive as well. This can be noted from observing that the operating region is much smaller than the temperature compensation region. Since each diode has its own characteristic, the only practical method of operating a laser diode in a widely temperature variant environment is to provide local temperature control, to within a few degrees C, and to regulate drive current by monitoring the optical power output of the diode. This most likely will require calibration of the drive network to match the characteristics of each specific diode.

Figure 40 illustrates the complexity required to operate a laser over the temperature range from 0 °C to +50 °C. To the best of our knowledge no one has demonstrated operation of a semiconductor laser diode over the full military temperature range. The circuit shown maintains a constant temperature of 0 °C ± a few degrees by monitoring the temperature of the laser diode chip which is mounted directly on a thermoelectric cooler and providing an appropriate control signal to the cooler. Electrical bias to the diode is controlled from an optical feedback tap which samples the output signal. This compensation network must also sample the digital drive signal and shut the laser off when no modulation is present.

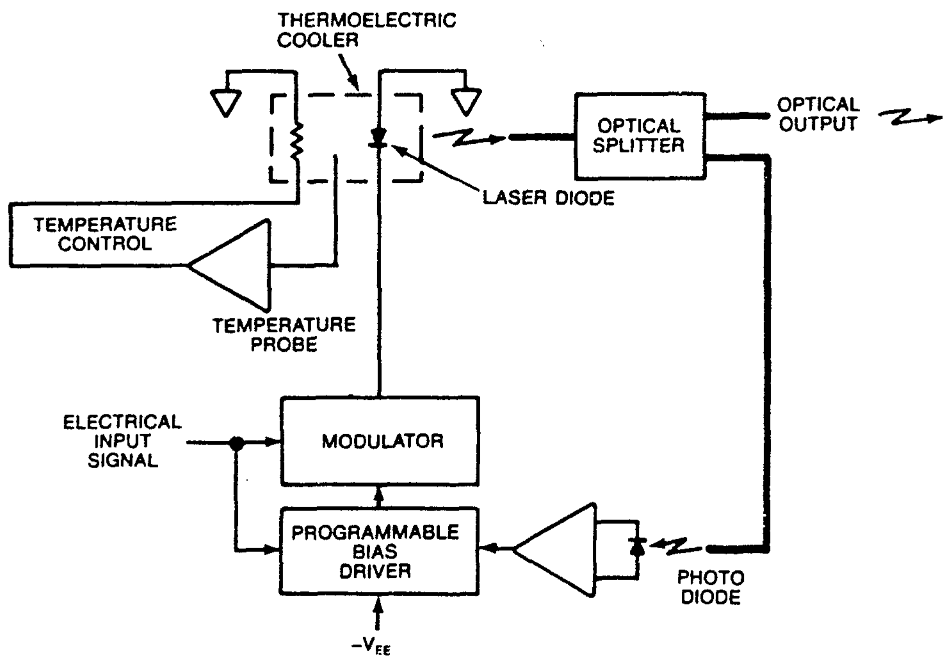

Figure 40. Typical Laser Drive Compensation Circuit

Another of the operational differences between LED sources and laser sources is their on-to-off operating region. Since the current through an LED can be turned completely off,

there is no optical power output from the device at logic 0 state. A laser, on the other hand, is generally current-biased just at the lasing threshold so that the signal drive current serves to pull the laser from the LED operating region (off state) to the lasing region (on state). Thus, even in the off state, the laser emits optical power into the system. This means that there is always optical power output from the transmitter, even when data is not being sent. This is required to maintain laser temperature and bias. Lasers are generally equipped with photodiodes to allow monitoring of their optical output. The signal from the photodiode is used in a feedback circuit to maintain laser output power at some average level. When no data is present, the stabilization circuit would tend to drive the laser so that the average output level is maintained. Thus, even with no data, the transmitter would generate an average signal level equivalent to the signal level when data is present. In data systems where the data is continuous, this situation does not cause a problem. In the HSDB, however, where data occurs in bursts the background noise (light) would degrade the S/N ratio of the network and must be eliminated by completely shutting off bias current to the diode. This would result in the need for a long preamble on each packet to allow the diode to be restabilized.

For the reasons noted above, it was decided that the LED represented the only reasonable choice for the optical source. It would be nice to have the additional power possible with a laser diode transmitter but the associated problems were felt to be too severe to justify its selection. Analysis in the other areas of concern showed insufficient rationale to reverse that decision, bandwidth being the only other characteristic in which the LED was inferior. This made it relatively simple to arrive at a peak power output specification. Available diodes were rated at -3 dBm maximum. This was reduced to -6 dBm as a result of tests performed on representative diodes to accommodate the temperature compensate circuit affect.

Optical Detectors

Two alternatives were studied to select the HSDB optical detector; these were the Positive-Intrinsic-Negative (PIN) photodiode and the avalanche photodiode (APD). Preliminary network design and selection of LED optical sources made it apparent that achieving a superior receiver design was key to success of the program. In keeping with the applications oriented objective of the analysis phase, the following device characteristics were studied:

a. Responsitivity
b. Bandwidth (risetime)
c. Ease of use
d. Reliability
e. Cost

As with the source analysis, the study was done assuming that little if any performance improvement could be expected during the initial application period. In effect, parts available at the time of study were judged to be representative of those available for use in early fielded systems.

Response and bandwidth are of principal importance in a detector since they directly drive the operating envelope of the HSDB. Table 8 summarizes the characteristics of importance typical for each type of detector.

Table 8. Photodiode Characteristics

DIODE	RESPONSITIVITY	RISETIME	BIAS	NEP
PIN	0.6 A/W	3 nS	-15 V	10^{-14} W/Hz
APD	75 A/W	2 nS	-300 V	10^{-12} W/Hz

Notable is the much higher responsivity of the APD. This is achieved because its avalanche mode of operation provides internal gain following the photon-to-electron conversion. Either diode has adequate bandwidth for use on the HSDB so a decision was made on the basis of sensitivity and ease of use.

Sensitivity refers to the minimum signal level (optical) required to produce a specific error rate. It is closely related to responsivity but also includes noise-equivalent power (NEP) and dark current. Since we had not defined the required S/N ratio at that time it was decided to compare the two approaches on the basis of minimum detectable signal (MDS). MDS is the point where signal = noise (S/N = 1).

For an APD:

$$\text{MDS} \approx 2/R \sqrt{q^B} \,(I_D + I_T)$$

For a PIN:

$$\text{MDS} \approx \frac{2q^B}{R} \frac{\sqrt{I_D + I_T}}{q^B}$$

Where
R = Responsitivity
B = Signal Bandwidth
I_D = Dark Current
I_T = Thermal Current
q = Electron Charge

Using typical values derived from component data sheets, MDS for an APD was estimated to be 8 X 10^{-10} Watts (-70 dBm) and for a PIN was estimated to be 10 X 10^{-8} Watts (-61 dBm). This shows an 11 dB advantage in sensitivity for the APD. While these sensitivities cannot be achieved in practical receiver designs, nevertheless the same relative difference in performance can be expected.

Ease-of-use characteristics favor selection of a PIN photodiode for two reasons. The first is bias voltage. In order to achieve avalanche gain the APD must be reverse biased by several hundred volts. Typical characteristics are shown in Figure 41. While the need for this bias voltage was not considered a terminal flaw, it did represent size, weight, and power penalties not assessed against a PIN diode receiver, even though an extra 10 dB of gain would be required of a PIN receiver. Temperature sensitivity was considered a more detrimental characteristic of APDs. Note from Figure 41 that the responsivity is highly dependent upon temperature. While some variation in gain is acceptable there would be no way to implement a receiver capable of operating over the temperature envelope expected of the HSDB without complex compensation electronics. This would probably include the need for temperature control similar to that described for laser diodes. In the end, it was decided that it would be impractical to operate APD receivers in a high performance military aircraft.

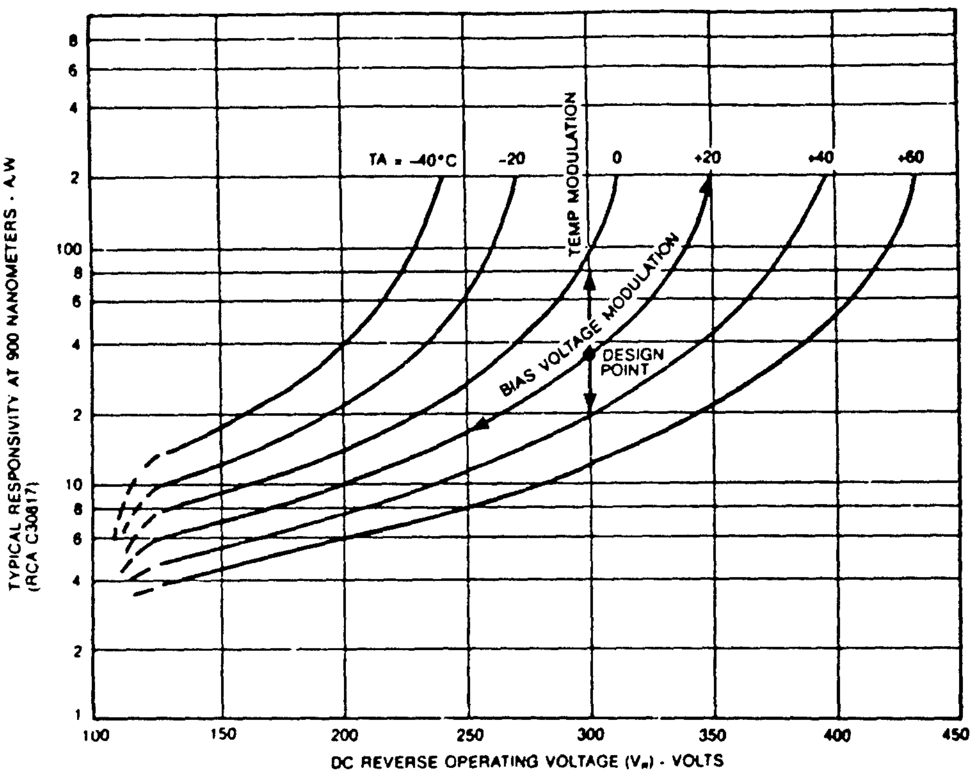

Figure 41. An Avalanche Photodiode Requires Controlled Bias and Temperature

For the reasons stated above, PIN diode detectors were selected for use on the program. The sensitivity and operating range specifications were determined by analysis and test of representative optical-to-electrical converter designs and PIN photodiodes from several sources. This indicated that -37 dBm was the theoretical limit of sensitivity. We adopted this figure as a design goal but derated the requirement to -35 dBm for system design analyses which followed. The ROR requirement was arrived at by assuming a worst-case 8-node network. Under worst-case conditions transmitter power is maximum (-3 dBm) and network loss is minimum (-9 dBm) requiring the receiver to operate with a -12 dBm signal applied. ROR then becomes 23 dB (-35 dBm to -12 dBm).

4.1.1.2 Topologies Analysis

The topologies analysis looked at various network configurations which could be used to interconnect 64 nodes in a broadcast network configuration. The three general topologies analyzed were: (1) Linear Bus, (2) Star Network, and (3) Hybrid Networks.

The analysis assumed a 30 dB optical power budget, the requirement which resulted from the power budget analysis. Three performance measures-of-merit were used to characterize each candidate topology. They were:

1. **Minimum Usable Signal (MUS)** - This parameter defines the required receiver sensitivity.
2. **Optical Signal Range (OSR)** - This parameter defines the maximum difference in an optical signal that the receiver needs to accommodate.
3. **Bus Dynamic Range (BDR)** - This parameter defines the maximum difference in optical signal level at any two parts of the network.

Additional criteria, including reliability, flexibility, expandability, radiation hardness, and installation limitations were also considered. The process involved four steps:

1. Define the alternative topologies
2. Define the measures of merit
3. Define the limits of each topology relative to number of terminals and data rate based on the maximum allowable bus loss for the minimum usable signal and the optical signal range and bus dynamic range
4. Select a single topology and characterize it

At the final step, the alternatives were traded against one another in order to select the topology to be developed during the remainder of the program.

For this initial first order analysis the best case performance for splices, connectors, and fiber was assumed. This was done to allow candidates to compete on the basis of topology considerations rather than have interconnect considerations cloud the issue. Interconnect considerations were analyzed during Step 4, after a single topology had been selected. Table 9 summarizes the loss equations of each of the candidates. Figure 42 graphs the function loss vs. number of terminals. Shown for comparison is the best possible theoretical case where the power available is split evenly between all receivers with no other losses. As can be seen, the only viable passive topology for 64 terminals is a single transmissive star; for this reason it was selected as the topology for the HSDB program. Rockwell believes, however, that the single star topology is not appropriate for use on high performance military aircraft and that work should continue to develop some form of linear bus topology.

Table 9. Topology Analysis Basic Loss Equations

APPROACH	COUPLING LOSS	THROUGHPUT LOSS
Linear Bus		
Unidirectional ON-TAP	.5 dB	.5 dB
Unidirectional OFF-TAP	10 dB	.5 dB
Bidirectional ON-TAP	3.5 dB	.5 dB
Bidirectional OFF-TAP	10 dB	.5 dB
Star Network		
Transmissive Coupler	10 log (N)	2 dB
Reflective Coupler	10 Log (2*N)	3 dB
Hybrid Network: appropriate selection from above (N = number of nodes)		

Analysis Of Linear Bus Topologies

Linear bus topologies may be either unidirectional or bidirectional and either passive or active by design. In a unidirectional tapped bus, directional couplers are used to tap the transmitters and receivers onto a single fiber. This is illustrated in Figure 43a. Typically a tap into the receiver can be accomplished with a 90/10 or 95/5 split providing 0.5-0.2 dB link throughput loss respectively and a 10 dB to 13 dB tap-off or reduction of the link power into the bus receiver. For tapping the transmitter into the bus, the throughput loss as well as the coupled transmitter power reduction is about 3 dB in commercially available couplers. Discussions with several coupler manufacturers, however, indicated that an asymmetrical coupler with two different fiber sizes could be fabricated with an estimated link throughput loss of less than 0.5 dB and a coupled transmitter power reduction of 0.5-1 dB.

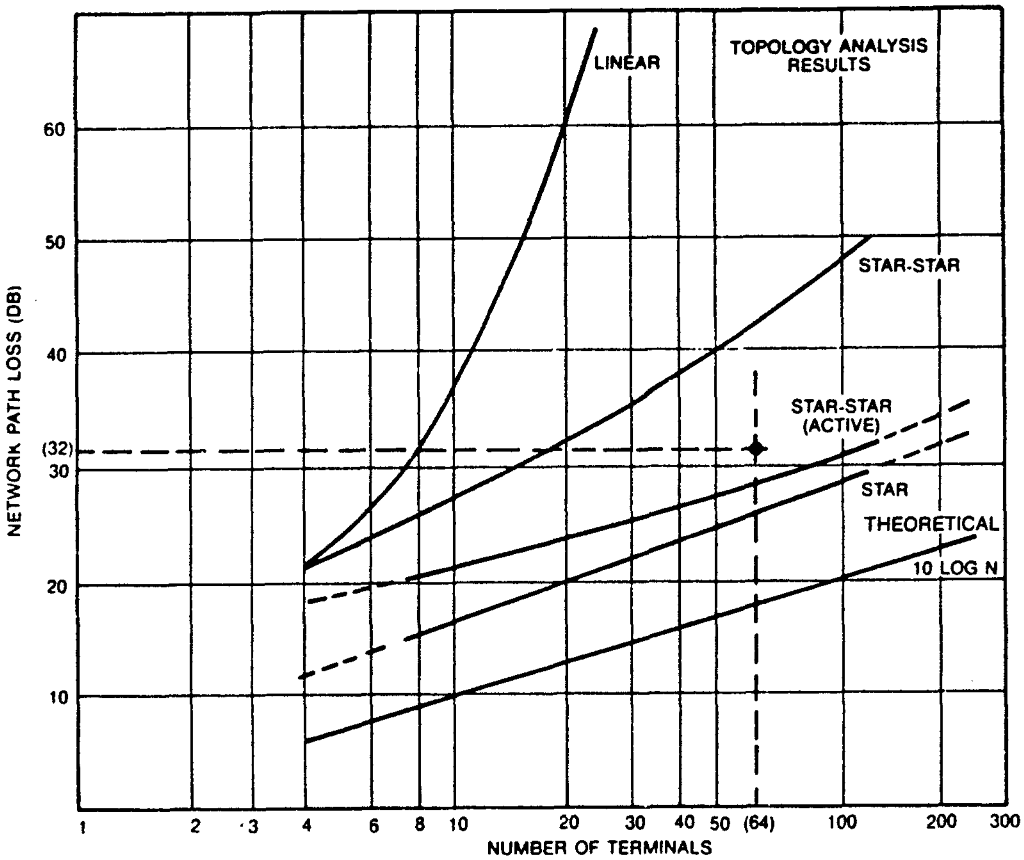

Figure 42. The Star Coupler Configuration is the Only Passive Candidate Meeting the System Loss Budget

A bidirectional tap can be accomplished several ways but each involves more loss than the unidirectional tap. Figure 43b shows one of the implementation alternatives. The theoretical Tx-link-Rx loss as well as the link throughput loss is shown in Table 10. These losses do not include the excess insertion loss of 0.5-1 dB of each coupler. Another technique, using separate "optimized" transmitter and receiver couplers similar to the unidirectional taps, appears to provide the best overall performance. This is due to the low link throughput loss.

A typical example of each form was analyzed. An example of a bidirectional passive linear bus is shown in Figure 44. The worst case loss from transmitter to receiver is shown in Table 11. As is evident from Table 11, this approach is not appropriate for a 64 node network since it greatly exceeds the allowed power budget. The maximum number of nodes for which this approach will work ir a practical installation is about 12. The actual number depends upon derating factors for service margins and interconnect loss, factors which were not addressed in this simplistic analysis.

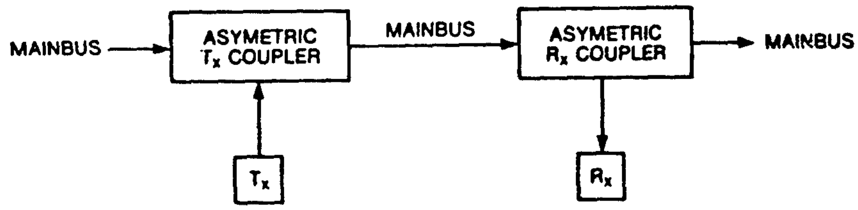

Figure 43a. Typical Unidirectional Coupler

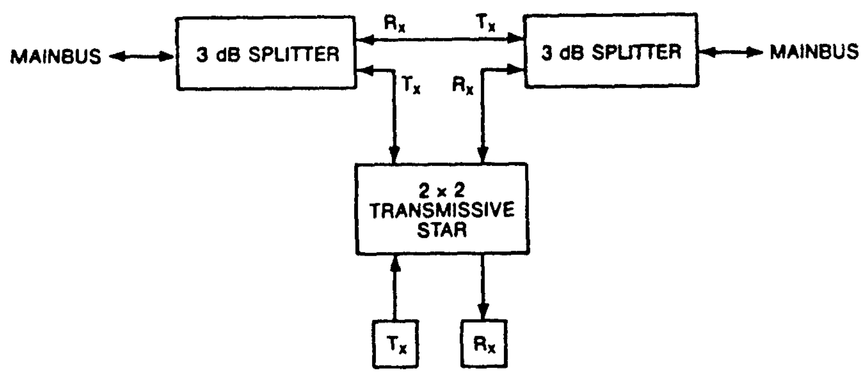

Figure 43b. Typical Bidirectional Coupler

Figure 43. Linear Bus Coupler Configurations

Table 10. Loss of Couplers

TYPE	Tx-LINK-Rx-LOSS	LINK THROUGHPUT LOSS
3 X 3 Transmissive Star	3+4.8+4.8 = 12.6 dB	4.8 dB
4 Port Reflective Star	6+6 = 12.0 dB	6.0 dB
3 dB Splitters	3+3+3+3 = 12.0 dB	6.0 dB
Combiner/Splitter	3+1+10 = 14.0 dB	0.5+0.5 = 1.0 dB

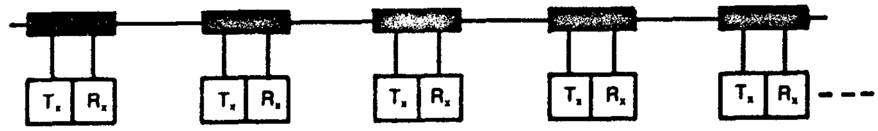

Figure 44. The Simplest Form of Linear Bus Topology Requires Bidirectional Couplers

Table 11. Network Loss Calculations for a Bidirectional Passive Bus

Maximum Loss	=	[ON-TAP Loss]
	+	[N x Transfer Loss]
	+	[OFF-TAP Loss]
For a 64 node network:		
Maximum Loss	=	3.5 dB + 62 x .5 dB + 62 x .5 dB + 10 dB
	=	75.5 dB

Repeaters could be inserted every 4 to 12 nodes, to resolve the loss problem. An almost unlimited number of nodes could be accommodated in this manner. Each repeater, however, adds additional cost and signal degradation to the network and also reduces its reliability. Ultimately, the combination of cost and risk associated with the repeater kept this approach from being accepted.

Another candidate topology, an unidirectional linear bus approach is illustrated in Figure 45. The worst-case loss for a 64-node network of this form is shown in Table 12. Note that the loss is somewhat less than that of the bidirectional linear bus approach but is still in excess of the power budget. It was not selected for the same reasons.

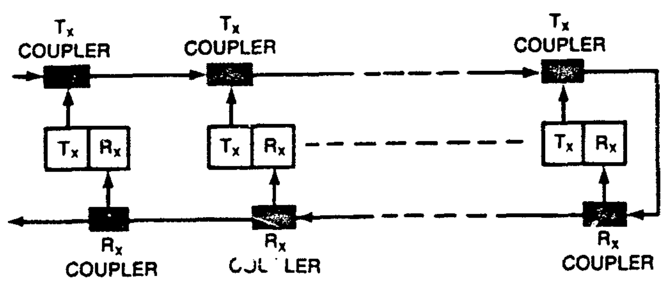

Figure 45. Unidirectional Passive Linear Bus Candidate

Analysis Of Star Network Topologies

Star networks may utilize single star or multiple star configurations, may utilize either transmissive or reflective star couplers, and may be either passive or have embedded gain

(active). As a precursor to the analysis itself the start-of-the-art of optical star couplers was extensively surveyed.

Table 12. Network Loss Calculations for a Unidirectional Passive Bus

Maximum Loss	=	[ON-TAP Loss]
	+	[2 x N x Transfer Loss]
	+	[OFF-TAP Loss]
For a 64 node network:		
Maximum Loss	=	.5 dB + 2 x 62 x .5 dB + 10 dB
	=	72.5 dB

Figure 46 shows two types of star couplers; (1) a transmissive and (2) a reflective. In the transmissive star, N ports are designated as input ports, and N ports are output ports. The optical energy on any input port is split more or less equally between all output ports. In a reflective star, the energy of any port is split between all ports and therefore any port may be designated as either an input or output port. For a symmetric configuration of N x N (input to output) ports a reflective star has fundamentally 3 dB higher coupling loss than a transmissive star, i.e., 10 log 2N vs. 10 log N. In addition to the coupling loss, star couplers have an insertion loss and a port-port variation (non-uniformity) each being in the range of 1-3 dB depending on the number of ports. This results in a total of 2-6 dB excess loss. Stars with up to 100 ports have been fabricated, however, for minimum cost and port-port variations, the practical limit of current technology is 64 ports. No testing had been done to establish reliability performance characteristics but several manufacturers expressed little concern in meeting all MIL-E-5400T [(7)], Class II environmental requirements given an opportunity to repackage the couplers for avionics applications.

The single passive transmissive star topology candidate is illustrated in Figure 47. It represents the simplest and most efficient topology possible since it operates as a simple power divider network. The simple loss equation of the network is:

$$\text{Loss} = 10 \text{ Log (N)} + \text{excess loss}$$

In practice excess loss is in the range from 1 dB to 3 dB, dependent upon N. For a 64 node network the expected loss, transmitter to receiver, is 21 dB. This is well within the planned power budget.

---

(7)MIL-E-5400T, "Military Specification, Electronic Equipment, Aerospace, General Specification For"

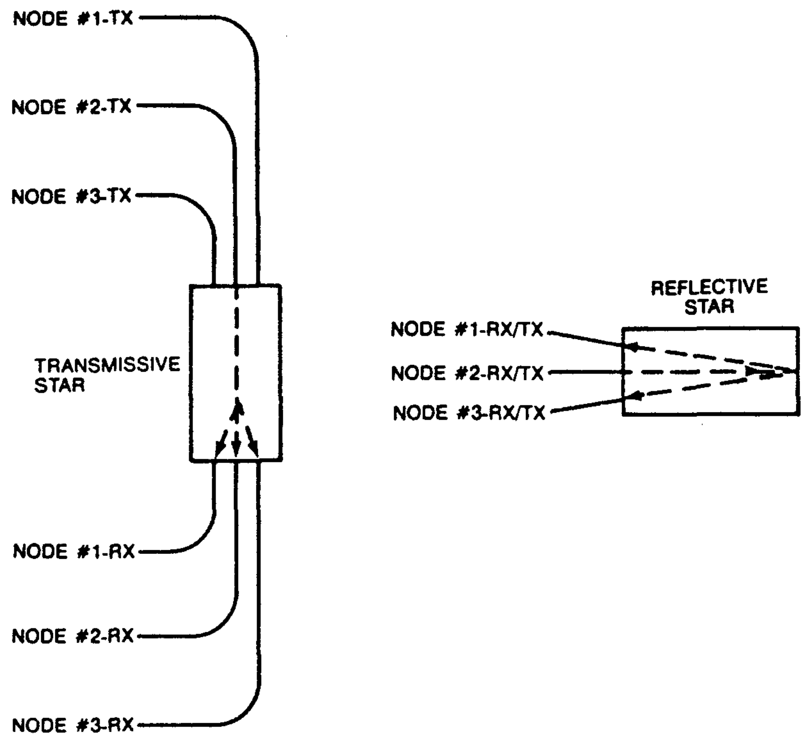

Figure 46. Two Configurations of Optical Star Couplers

A single passive reflective star topology would exhibit similar characteristics, with approximately 6 dB higher loss. The only other characteristic different from the transmissive star network is that only half the number of fibers is required. The positive implications of this are in improved installation characteristics in an aircraft. Reflective stars cannot be used in a true broadcast topology, however, because the transmitting node cannot receive its own signal through the network. This may complicate built-in test (BIT) design and also may restrict protocol options. Principally for the latter reasons a reflective star approach was rejected for this program.

The principal disadvantage of either single star topology is that the cables from all TRU must be run to the central coupler. In an aircraft, ship, or submarine this increases the initial installation cost due to the increased number of bulkhead connectors required. In addition, there is little flexibility for adding new terminals at arbitrary locations. One solution to this is to provide a distributed bus topology such as a star-star as shown in Figure 48 or hybrid linear-star or star-linear topology which will be discussed later. The network loss for various numbers of terminals in a quad-cluster star-star topology is shown in Figure 49 and in Table 13. The performance of this topology can be easily improved by adding a single repeater (or two for

redundancy) at the central star. This approach was not selected for this program, however, because it would increase program cost and risk with little advantage to offset these.

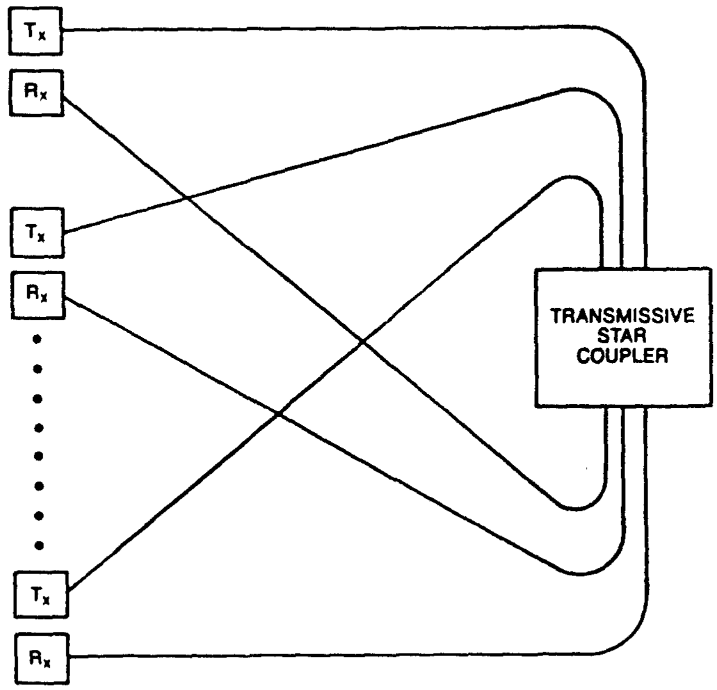

Figure 47. The Single Transmissive Star Topology Minimizes Loss

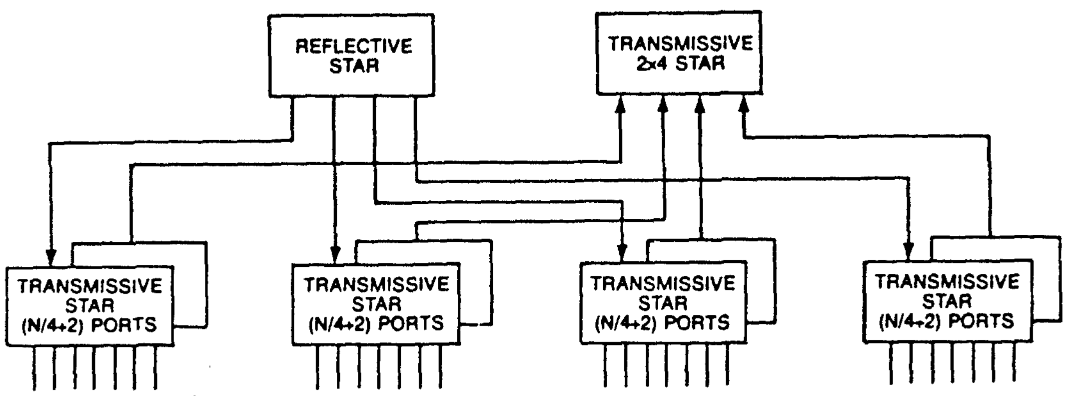

Figure 48. This Active Star-Star Topology Features Redundancy

Hybrid Topologies

Two hybrid topologies combining stars with a linear bus concept were investigated because they provided four separate node clusters with the potential of improved performance over a simple linear bus. The first, a star-loop, is shown in Figure 50. This topology could be

made active with a redundant repeater at the star similar to the active star-star. The second is a loop-star as shown in Figure 51.

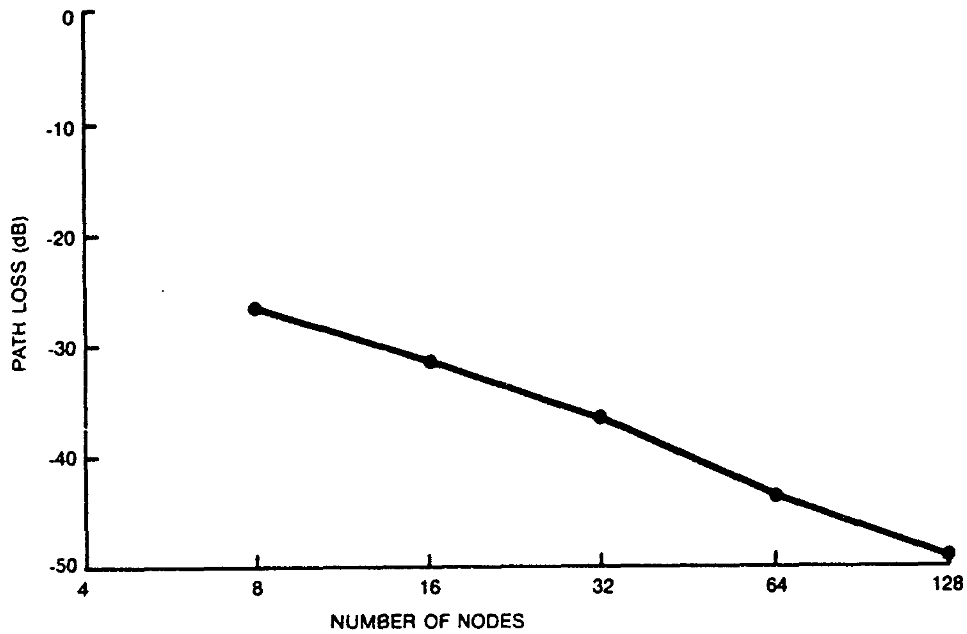

Figure 49. Quad-Cluster Star-Star Topology Network Loss

Table 13. Quad Star-Star Network Loss Equations

Maximum Loss	=	Tx Star Concentrator Loss (10 Log (N) + 2 dB)
	+	Tx Central Star (10 Log (4) + 2 dB)
	+	Rx Central Star (10 Log (8) + 3 dB)
	+	Rx Star Splitter (10 Log (N) + 2 dB)
	=	20 Log (N) + 24 dB
For a 64 node network:		
Maximum Loss	=	60 dB

Initial analysis of these revealed very little reduction in bus loss over a simple linear loop and therefore a detailed analysis was not performed.

4.1.1.3 Optical Interconnect Analysis

Several considerations are involved in evaluating the optical cabling for a fiber optic data bus. They include:

a. Fiber type
b. Cable type and construction
c. Optical connectors
d. Optical splices

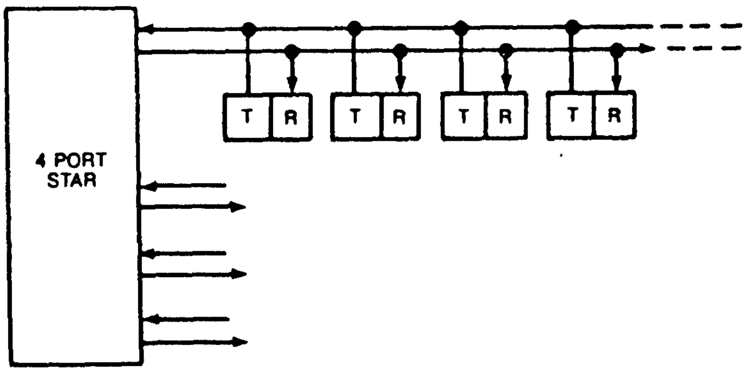

Figure 50. The Star-Loop is a Hybrid of Linear Bus and Star Topologies

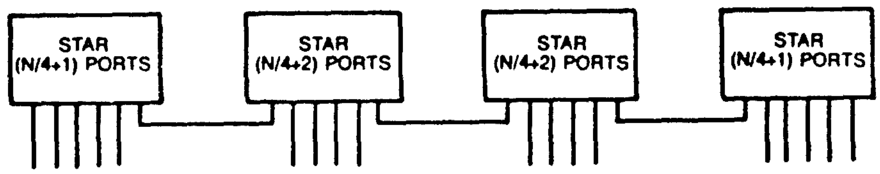

Figure 51. The Loop-Star Topology Supports Isolated Clusters of Nodes with Minimum Interconnect

Another consideration in the interconnect analysis was reflections. Reflections result from an index of refraction discontinuity at connectors, poor splices, or mismatched fiber types. For example, with a star coupler, the main signal passes through the link, however, part of the signal is first reflected at the input connector (8%) and then again at the output connector (8%). The resulting reflected signal is down 22 dB with respect to the main signal and delayed by 1 μS (1 nS/meter). This reflected signal becomes a problem if it overlaps the next bus transmission and shows up as noise superimposed on the data. The following consideration which must be given to minimize reflections include:

a. Use splices rather than connectors

b. Use index matched (wet) connectors

c. Optimize receiver sensitivity so as to prevent detection of the reflections.

Optical Fiber

A major effect to consider in selecting the optical fiber type is modal noise. All multimode optical fiber transmission systems suffer from modal noise to some degree. The origin of the effect is similar to radio multipath fading and distortion and is the result of interference between different signal propagation paths. In fiber, these paths are the fiber (transverse) modes, and the interference is evident by a speckle pattern, i.e., uneven spatial distribution of light intensity exiting the fiber end. While the total pattern intensity is predictable, it is not possible to describe exactly how the energy is distributed across the fiber cross section thus when some portion of the power is lost by cross section selective loss, such as misalignment of a fiber joint or coupling within a star coupler, an uncertain loss occurs. Source frequency instability and mechanical disturbance of the fiber can alter the speckle pattern and therefore the loss may vary from moment to moment. The resulting received power variation with time or with source modulation is known as modal noise, or modal distortion.

Guidelines for minimizing modal noise include increasing the number of propagating modes by:

a. Using large core fiber
b. Using high numerical aperture fiber
c. Operating at short wavelength
d. Using broadband sources

Table 14 shows the fiber waveguide tradeoffs. The 100/140 μm step index fiber is marginally capable of providing the required bandwidth; 50/125 μm graded index fiber has relatively poor modal noise characteristics and couples poorly. The most serious limitation of the 100/140 graded index fiber is that it was a relatively new design (in 1984) and was manufactured only by Corning and Valtec. Since only a small quantity of cable was required for the proposed program, supply was deemed to not be a problem. Over the long term, it appeared that the 100/140 graded-index fiber would replace the current 100/140 step-index fiber for the commercial data transmission industry. For those reasons, the 100/140 μm graded-index fiber operating at 850 ηm wavelength was selected. To summarize:

a. Its large core, high NA, and operating wavelength will support a large number of propogating modes, thus minimizing losses in connectors.

b. Its large core enables greater LED-coupled power, thus extending the application of LEDs.

c. The core-clad geometry makes it easier to make low excess loss star couplers.

Optical Cable

Optical fibers are rarely used in their native form because of their delicacy and because they are not easily connectorized. Rather, they are provided as cable assemblies of one or more fibers in a protective jacket of some kind. Three basic types of optical cable can be present in an optical HSDB, depending on the topology. They are:

1. Single fiber cable-linear bidirectional bus
2. Two fiber cable-linear loop bus and T/R interconnection to a star coupler.
3. Multi-fiber (bulk) cable-interconnections from a star to concentrations of TRU.

Use of the latter may be desirable to minimize bulkhead penetrations. The size and weight of optical cables are significantly less than wire cables or coax; strength and environmental performance of current cable technology will meet anticipated requirements. If nuclear radiation requirements are imposed special optical fiber and cable construction must be employed to improve survivability. In any event it was determined that the technology required to provide almost any form of optical cable was already in existance and so no further study was done in this area.

Table 14. Fiber Waveguide Tradeoffs

	Fiber Geometry		
CHARACTERISTIC	50/125 μm GRADED-INDEX	100/140 μm STEP-INDEX	100/140 μm GRADED-INDEX
Numerical Aperature	.20 - .22	.28 - .30	.28 - .30
Number of Modes	Low	High	High
Bandwidth	800 MHz-km	50 MHz-km	200 MHz-km
Attenuation at 850nM	3-4 dB/km	5-6 dB/km	4-5 dB/km
Wavelength	850-1300 nm	850-1300 nm	850-1300 nm
Coupler Capability	Poor	Good	Good
Availability	Good	Good	Limited
Modal Noise Reduction	Poor	Good	Good

Optical Connectors and Splices

Fiber optic connectors offer the most convenient method of interconnecting different parts of the network. Connectors in a fiber optic network do not exhibit the benign characteristics of electrical connectors in the companion coaxial HSDB. Optical connectors which are suitable for use in a HSDB are of two basic types: (1) single fiber and (2) multi-fiber. The single fiber connectors are low cost, easily installed, and typically offer lower loss than multi-fiber connectors. The connector loss depends on the fiber size as well as the quality (and cost) of the connector. For 100/140 μm fiber, losses vary from 0.5 to 1.5 dB depending on connector quality. Available multi-fiber connectors have the advantage of simplifying a bulkhead penetration and provide quicker connect/disconnect of a multi-fiber cable. Although there is no fundamental reason for higher loss in a multi-fiber connector, the losses in currently available connectors average approximately 0.5-1 dB more than the loss in a single fiber connector. Maturity resultant from passage of time may eventually eliminate this differential. Since connectors were not a major focus for this program the decision was made to use single-fiber connectors to minimize the loss. This is not necessarily a good choice for actual aircraft installations.

Splices are a low-loss alternative to connectors where the connection can be permanent. Most laboratory and factory fiber splices are performed using a flame fusion technique. For field installation, maintenance and repair, elastomeric splicing has been identified as the best currently available technique. The specifications, features and benefits of this splice are shown below.

- Designed to splice 50/125 μm or 100/140 μm fiber
- Strain relief is provided for .9-1.0 mm tight buffered fiber
- Splice material: Polyester Elastomer
- Splice Housing material: Polyester
- Overall Dimensions: 3.75 in. x .40 in. x .375 in.

4.1.2 Synthesis Phase

Results from the analysis phase provided baseline design guidance for the synthesis (design and development) phase of the program. This guidance is summarized below:

a. Topology: Single 64-port transmissive star
b. Transmitter: LED
c. Receiver: PIN photodiode
d. Power Budget: 30 dB

The synthesis phase began by flowing-down this high level design and the generalized network requirements and resulted in a detailed design for each element of the fiber optic HSDB network. This was accomplished in the following sequence:

1. The loss budget for the network was determined. This included allocating loss to each network element in the path from transmitter output port to receiver input port.
2. The power output requirement for the transmitter was determined by selection of a specific LED type to be used for the program.
3. Receiver sensitivity and operating range requirements were determined from the transmitter power output requirement and minimum/maximum network configuration.
4. Network modulation format was selected to be compatible with the transmitter and receiver requirements.

In effect, the synthesis phase began by reaffirming the generalized requirements from the analysis phase. This process resulted in the requirements set which are summarized in Table 5.

Development specifications were prepared for the TRU, then design of hardware was begun. The specifications shown in Table 5 were modified to some extent during the course of the program as a better understanding of the technology was achieved.

Step 1: Determining Network Loss Requirement

Determining the loss budget for the HSDB network consisted of identifying each element exhibiting a loss characteristic between the transmitter output port and the receiver input port. For this purpose, the synthesized aircraft installation shown in Figure 52 was adopted. Notice that the minimum configuration contains 8 nodes, 2 connectors and a small length of fiber, the maximum path includes 64 nodes, 4 connectors and 100 meters of fiber.

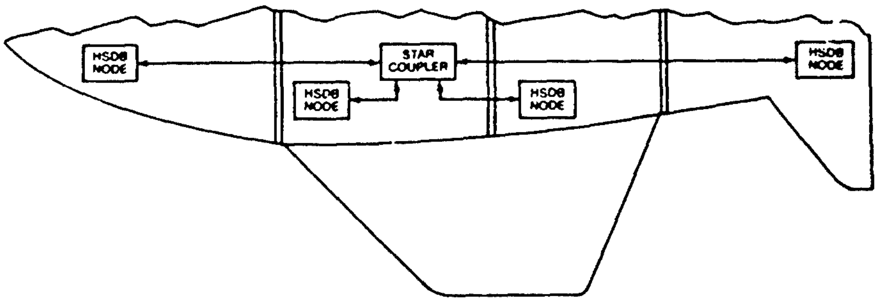

Figure 52. This Aircraft Installation was Assumed for Determining the Network Loss Requirement

The loss budget for the minimum configuration is:

Coupler loss (8 nodes)	9 dB
Coupler excess loss	0 dB
Connector loss (2 at 0 dB)	0 dB
Fiber loss	0 dB
Total Loss	9 dB

The loss budget for the maximum path is:

Coupler loss (64 nodes)	18 dB
Coupler excess loss	3 dB
Connector loss (4 at 1.5 dB)	6 dB
Fiber loss	2 dB
Total Loss	29 dB

Note that no single network will exhibit this loss range. This range encompasses the network envelope between 8 nodes through 64 nodes. Note also that this budget does not accommodate several losses which must be considered for production aircraft installations. These include:

a. Temperature affects on fiber, couplers and connectors.

b. Manufacturing and lifetime variations for sources, sensors, and other elements of the network.

c. Nuclear affects on sources, sensors, fiber and couplers.

The requirement was selected as being appropriate for the laboratory demonstration which was part of this program. It is not, however, adequate for production aircraft; that level of design was outside the scope of this program. Rockwell did not consider this to be a serious deficiency, however. Technologies developed by this program were focused in the areas of receiver and transmitter design, areas equally important no matter what topology is eventually chosen for production installations.

Step 2: Determining The Transmitter Power Output Requirement

The transmitter power output requirement was set to "all that we can get" as a design goal. What this really meant was to determine a requirement which could be met using production LEDs, when operated over the full MIL-E-5400T, Class II environment (-54 ^{o}C to +95 ^{o}C).

The process of selecting a LED turned out to be quite simple. At the time there were just two alternatives which could be procured off the shelf, which had an advertised output power of

-6 dBm or above and rise time less than 5 nS: (1) Laser Diode Labs LDT-474 and (2) the Plessey CLX-045. Samples of each were procured and tested. The results of the evaluation test are shown in Table 15.

Table 15. LED Qualification Test Results

CHARACTERISTIC	LDT-474	CLX-045
Power Output	Adequate	Adequate
Risetime	Marginal	Adequate
Temperature	Failed	Passed

From this screening exercise, the Plessey CLX-045 LED was selected for use on the program. The next step was to characterize its performance over temperature so that the power output requirement could be set. Figure 53 shows the average temperature vs. power output function for a typical LED. Since the LED could be safely operated to 150 mA bias to compensate for the drop in optical power produced at higher temperatures, -6 dBm peak was selected to be the power output requirement for the transmitter.

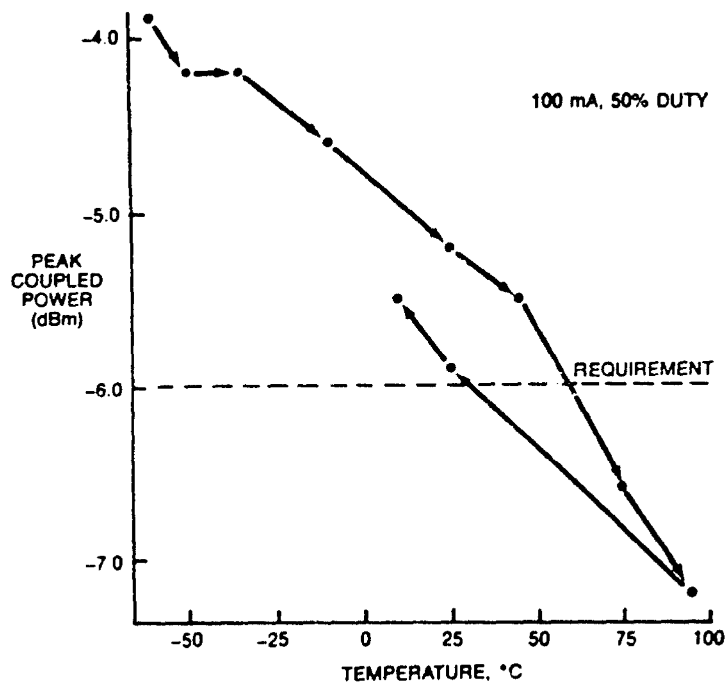

Figure 53. LED Coupled Power vs. Temperature

Step 3: Determining Receiver Operating Range and Dynamic Range Requirement

The receiver operating range (ROR) requirement could be determined after establishing the transmitter power output requirement and the network loss budget.

Maximum Signal Operating Point:		
Transmitter Power Output	-6 dBm	peak
Maximum Network Loss	29 dBm	
Receiver Min Signal	-35 dBm	peak
Maximum Signal Operating Point:		
Transmitter Power Output	-6 dBm	peak
Minimum Network loss	9 dBm	
Receiver Max Signal	-15 dBm	peak

Note that no receiver will see this variation in signal power since minimum network loss (8 nodes) and maximum network loss (64 nodes) are not consistent. Rockwell adopted this design requirement so that the same receiver design would operate reliably in either a minimum or a maximum size network without the need for adding attenuation to the interconnect.

Receiver dynamic range refers to the difference in signal level which the receiver must accommodate in a relatively short time such as between successive transmissions. In the case of the HSDB, it is the difference in signal level between two adjacent packets. The dynamic range requirement while operating in a network with a single star topology is minimal. The receiver must only accommodate variation in loss in the order of 8 dB. Since Rockwell felt that the single star topology was not appropriate for production aircraft, the decision was made to set the dynamic range requirement at a point where linear bus topology could be accommodated. In the absence of a specific design for a linear bus topology, it was decided to get the dynamic range requirement equal to the ROR requirement, 21 dB. This appeared appropriate from the topology trade study results.

Step 4: Line Code Transmission Format

Determining the requirement for line coding and format for the HSDB involved selecting from several more or less standard alternatives which had proven to be acceptable for fiber optic packet networks. The alternatives surveyed were:

a. On-Off Equivalent Manchester
b. Block Code (4B5B, 8B10B, etc.)
c. Miller

Manchester is one of the simplest line codes. It provides high clock content and error monitoring in a binary code. A disadvantage of Manchester is the need for twice the transmission bandwidth of NRZ transmission, and logical processing at twice the bit rate is required. The additional transmission bandwidth is readily available in optical fiber but available components and circuit technology make bandwidth a concern.

Block codes and Miller are formats which require coding/decoding logic operating at clock rates at or close to the data rate. For start and end of message delimiters, it is necessary to send unique codes which do not appear within the data, and so the code must replace these data sequences with alternative sequences using a block or bit substitution. Bit substitution is one of the simplest logic encoded codes to implement. Phase shift encoded differential biphase is encoded from a bit rate clock, and encoding and decoding use RF building blocks, e.g., balanced mixers, power splitters which are readily available to 300 MHz and greater.

Step 4 resulted in Manchester format begin chosen. Manchester allowed short preambles for receiver signal acquisition and generally simplified the design of the receiver. The LED already selected had adequate bandwidth to accommodate the 2X baud rate required. An 8-bit preamble and 4-bit-time invalid Manchester start delimiter and end delimiter were specified.

4.2 Transmitter Design

For relatively low rate transmission using LEDs, little difficulty exists in designing a transmitter circuit. Data modulation may be DC-coupled through to the LED and any data format or message length may be accommodated. For rates much higher than 20 Mbps, careful high frequency design is required and the logic must be designated using ECL. LEDs which have fast risetimes must be selected. Speed is traded off against optical power (which is lower for high-speed LEDs), to a maximum modulation rate of 100 MHz. Figure 54 shows the transmitter circuit designed for the fiber optic HSDB. The transmitter provides compensation to hold the peak optical power output constant over the operating temperature range. This is easily accomplished by controlling bias current to the LED switching transistor. Figure 55 shows effect of the compensation circuit designed for the HSDB transmitter. It holds peak power output constant within 0.2 dB over the entire operating temperature range.

The transmitter must also provide a monitor function and override control to extinguish the LED should an error occur in the drive circuit. This is not shown in Figure 54 but was included in the prototype hardware.

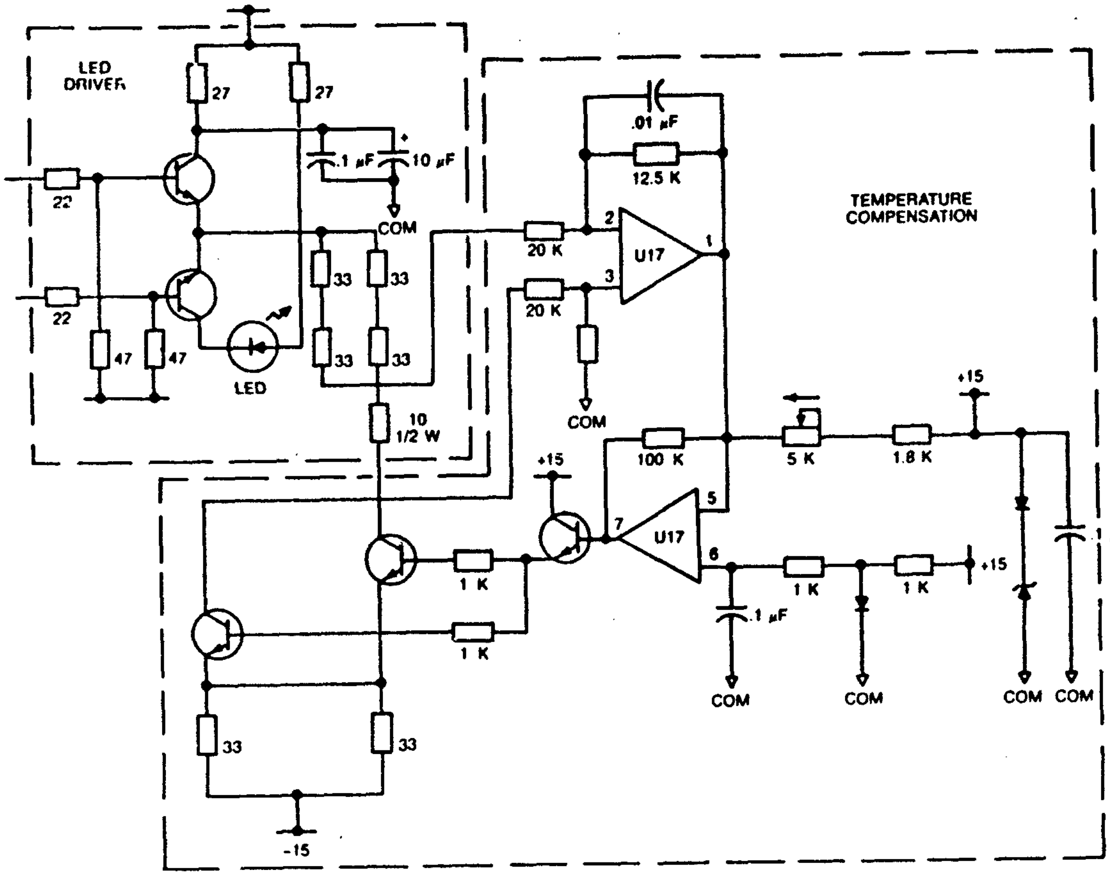

Figure 54. Fiber Optic Transmitter Simplified Schematic Diagram

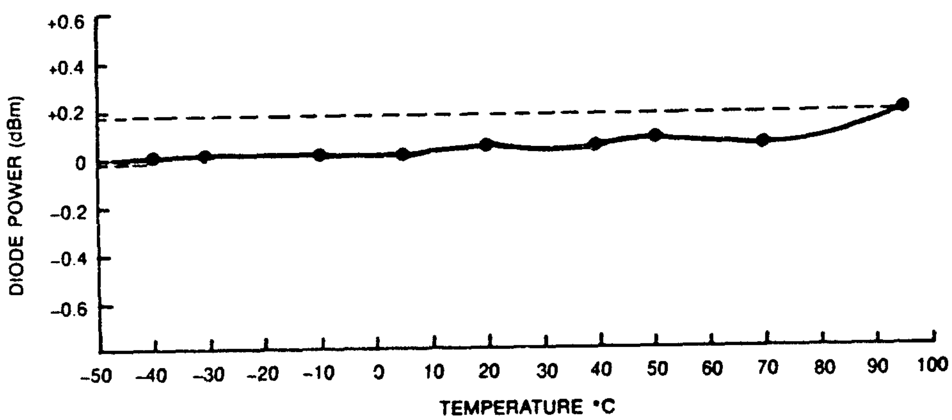

Figure 55. Transmitter Output Power After Compensation

4.3 Receiver Design

Development of a superior receiver design was the principal focus of design activities during Task II. The receiver for the HSDB network needed state-of-the-art sensitivity and dynamic range capability in order to meet the requirement.

The mean level shifting of unipolar burst transmission imposes strong restraints on receiver design if good bus efficiency is to be maintained. The three basic approaches to data detection are shown in Figure 56. Reasonably fast response to changing signal levels can be obtained by conventional ac-coupled receiver design with minimum coupling and AGC time constants, but a significant efficiency penalty remains at low data rates and data coding is restrained to avoid loss of frequency content. Alternately, the automatic gain control (AGC) time constant may be switched so that stabilization on the preamble is rapid, and then a slow time constant is used during the data transmission. Receivers, which do not have AGC, use an adaptive circuit to select a data decision threshold at nominally half the height of the signal peak, based on a rapid assessment of the signal amplitude during the preamble.

All of the above receiver schemes suffer a delay before valid data can be extracted. The preamble length must be increased to ensure that sufficient valid preamble remains for clock synchronism. The two latter schemes require high speed switching of analog signals at the appropriate times during reception, and the last receiver type has a more limited dynamic range since it has no AGC.

High bit rate reception may also be handled efficiently with a short time constant ac coupled receiver when the signal is any biphase code or other reduced low frequency content code. Coupling capacity time constants become small compared to the fixed bus inter-message dead time resulting from propagation delays.

No AGC and differential biphase code and RF demodulation components is used on the HSDB receiver shown in Figure 57. It consists of an optical to electrical converter with gain, a DC restoration filter, a digital comparator, a clock recovery circuit, and a Manchester demodulator.

Optical to Electrical Converter

The optical-to-electrical converter is the most critical function of the fiber optic receiver. It sets the performance in the areas of sensitivity, operating range and dynamic range. Figure 58 shows the design of this function. Light pulses comprising the signal are focused on the junction of a PIN photodiode biased to operate in its photoconductive mode. Photons striking the detector result in a proportional electrical current from the diode. This electrical signal is amplified by a high-gain, low-noise preamplifier. The preamplifier is based on a well known

transimpedance design first published by Ogawa and Chinnock. [8] While the design is relatively straight-forward, fabrication of prototypes proved to be troublesome. The extremely high gain and high input impedance required careful layout and shielding. This was successfully accomplished on the third printed circuit board layout.

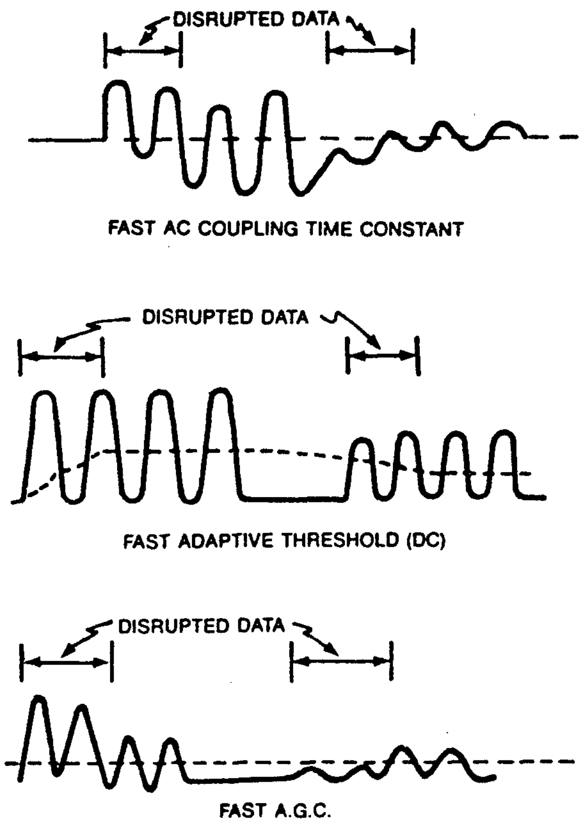

Figure 56. Three Different Approaches to Receiver Coupling

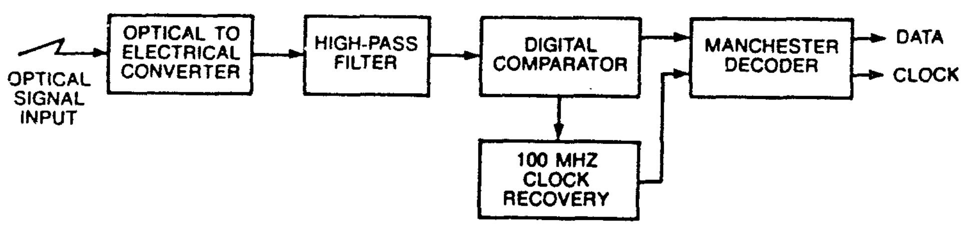

Figure 57. Fiber Optic HSDB Receiver Functional Block Diagram

(8)K. Ogawa and E.L. Chinnock, "GaAs Transimpedance Front-End Design For A Wideband Optical Receiver," Electronics Letters, 1979

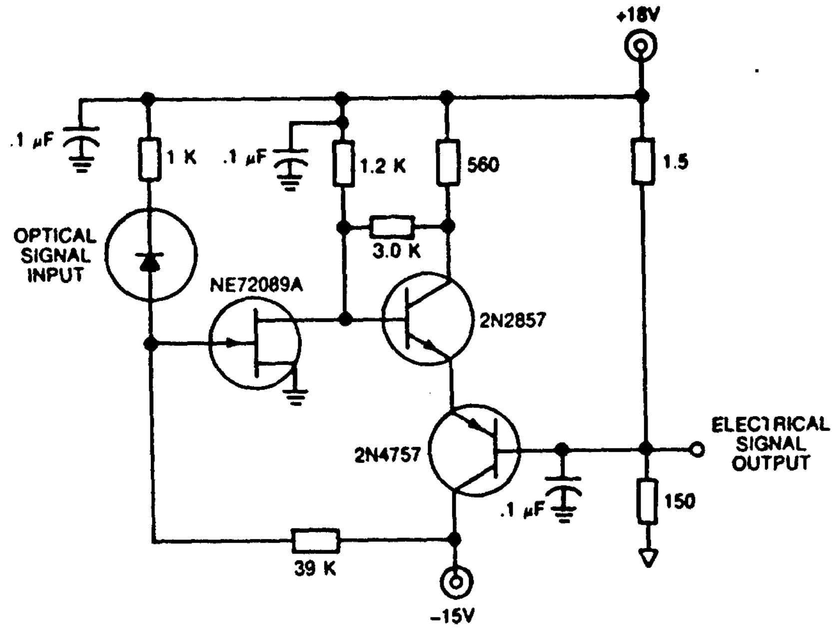

Figure 58. A Transimpedance Preamplifier Provides Low-Noise Amplification of the Photocurrent

High Pass Filter

The preamplifier is ac-coupled to following receiver stages using a high pass filter. Although DC-coupling offers the advantages of simplicity and added performance, a totally direct coupled optical receiver is not practical because of high temperature leakage current associated with the photodetector and the input FET of the preamplifier. AC-coupling, therefore, is required between the preamplifier and post amplifier. Although this solves the DC drift problem, care must be taken in the selection of the coupling time constant so that receiver acquisition time or signal droop is not unduly degraded.

In a receiver designed for the PAVE PILLAR HSDB, signal droop is not a problem since transitions occur during every data bit time. The only exception is during the start and end message delimiters where the signal remains constant for 1.5 data bit times. If one assumes a droop of 40% of initial value is acceptable, a time constant can be calculated by the following expression:

$$e^{-\frac{NT}{RC}} = 0.4$$

where RC is coupling time constant, T is the bit time and N is the number of bits. If N = 1.5 which is its maximum value, and T=20 ηS, RC is found to be 3.28e^{-8}. Receiver stabilization time can be similarly calculated. Because of the very large intermessage gap, receiver operating points

will have to be completely stabilized before every reception. Stabilization time, therefore, is the time it takes the coupled unipolar optical signal to drift below the baseline so that the data detector can reliably determine signal polarity. If the same criterion is used as before, it is assumed that a reliable signal is obtained when the signal has drifted to 40% of its final value. They may be expressed as:

$$e^{-\frac{NT}{RC}} = 0.6$$

where the variables are the same as before. Using the value of RC calculated for acceptable droop, N is found to equal 0.84 bit times. This means that the minimum droop requirement will be met for amplifier stabilization times of less than one bit time. Achieving short receiver amplifier stabilization times, therefore, was not difficult.

Clock Recovery Design

When data is received in a serial bit stream, a timing reference is required to decode the data. The reference should be a clock at the bit rate, phase synchronous with the data. With a noise-free signal, the clock phase may be determined from any data transition, but an integration of several data transition times is required to obtain averaged timing from data with noise. Averaging must be done before valid data is transmitted, and a preamble having a high transition density is generally transmitted for this purpose.

Clock recovery schemes, which would be suitable if NRZ or low-transition-density coded data were selected, would generally require mode switching to be able to acquire clock very rapidly and then to maintain clock with adequate stability. Mode switching techniques include a phase-locked loop with a high-speed acquisition mode, and a multiphase crystal clock with clock phase selection. The latter techniques would offer greatest stability and fastest acquisition time if parallel paths are used in the phase selection process.

The simplest clock acquisition circuit, and the approach chosen for the HSDB receiver, comprises an electrically resonant circuit tuned to the bit rate, which resonates on energy extracted from the data stream. The required phase averaging is achieved by choice of resonant Q. Resonant clock recovery is ideal for data transmitted with high levels of clock content throughout the message, e.g., any biphase-encoded data, but it is difficult to implement over a wide temperature range. Substantial engineering was required to arrive at a design which operated reliably over the desired temperature range.

4.4 Proof Of Concept Testing

The design of the fiber optic HSDB TRU was validated using various breadboard and brassboard test configurations. All critical circuits were breadboarded and tested over temperature prior to being integrated into breadboard TRU. Each breadboard TRU consisted of two assemblies, (1) a transmitter circuit board and (2) a receiver circuit board. Two of each were built and tested in a simulated fiber optic HSDB network. Finally, 6 TRU of brassboard configuration were fabricated, tested, and characterized. The test and characterization equipment developed for this purpose is described in Section 6.0 of this report. Figure 59 is a photograph of the brassboard receiver circuit board; Figure 60 is a photograph of the brassboard transmitter circuit board.

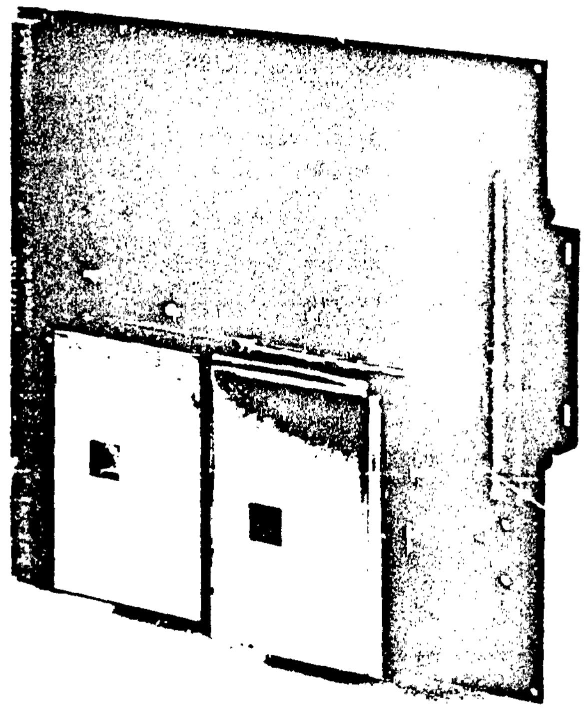

Figure 59. Brassboard Fiber Optic Receiver Circuit Card

Characterization

Characterization is defined as the process by which parametric performance of the design is proven. The following characterization testing was performed on the six brassboard transmitters and receivers:

a. Transmitter output power
b. Transmitter output waveform
c. Transmitter clock stability

d. Transmitter synchronization waveform
e. Transmitter timeout override
f. Transmitter switch waveform
g. Transmitter output noise
h. Transmitter modulation
i. Receiver acquisition range
j. Preamble response time
k. Receiver dynamic range
m. Receiver input impedance
n. Bit-error rate

Performance of the transmitter/receiver units was proven over the temperature range of -54 °C to +95 °C.

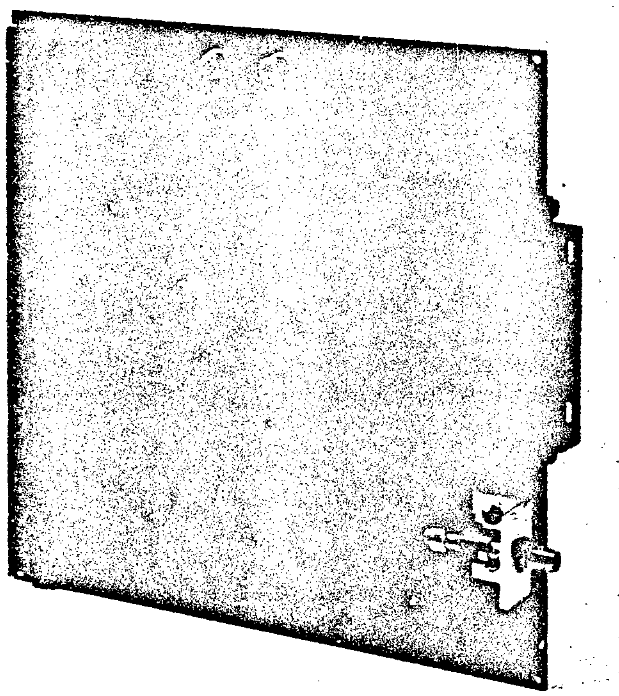

Figure 60. Brassboard Fiber Optic Transmitter Circuit Card

Demonstration

A demonstration of fiber optic network HSDB technology was conducted using three brassboard TRUs connected into the HSDB system demonstration equipment. This demonstration showed three HSDB terminals operating in a simulated HSDB network. The network was emulated using a fiber optic network emulator panel shown in Figure 61 and also using a 64-port star coupler. Figure 62 shows the HSDB demonstration equipment as configured

for Task II ATR demonstration. The Task II demonstration was accomplished with the same test equipment as used for the Task I demonstration except for substitution of the fiber optic network emulator. Figure 63 shows a typical waveform monitored on the network.

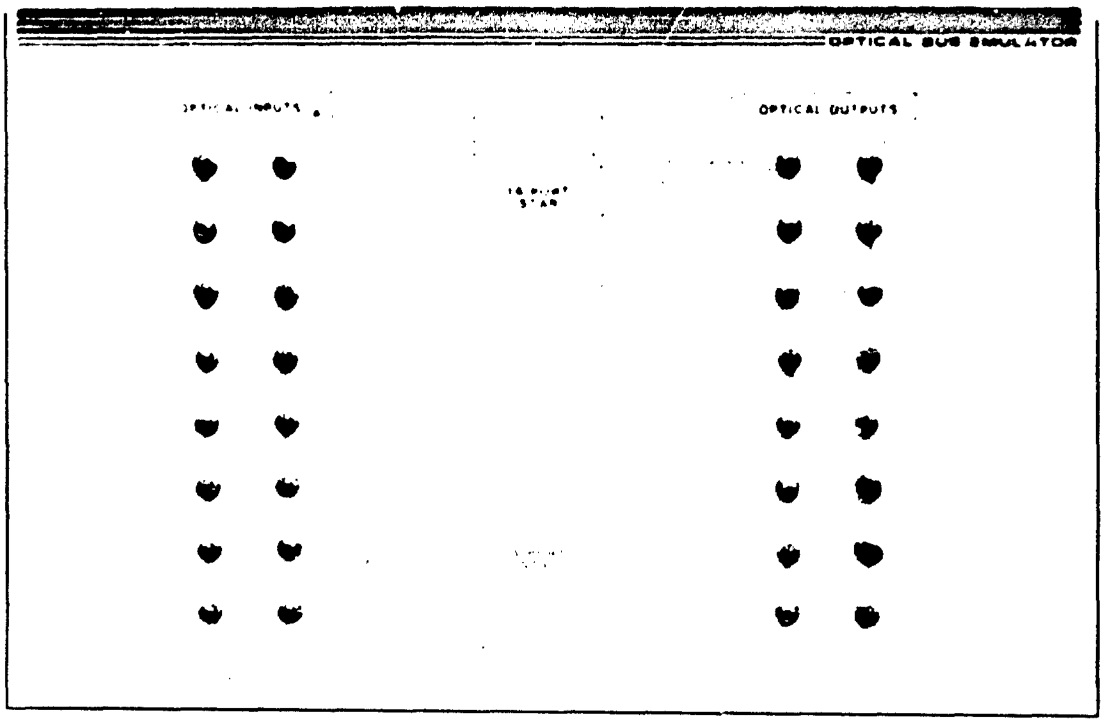

Figure 61. Fiber Optic HSDB Network Emulator

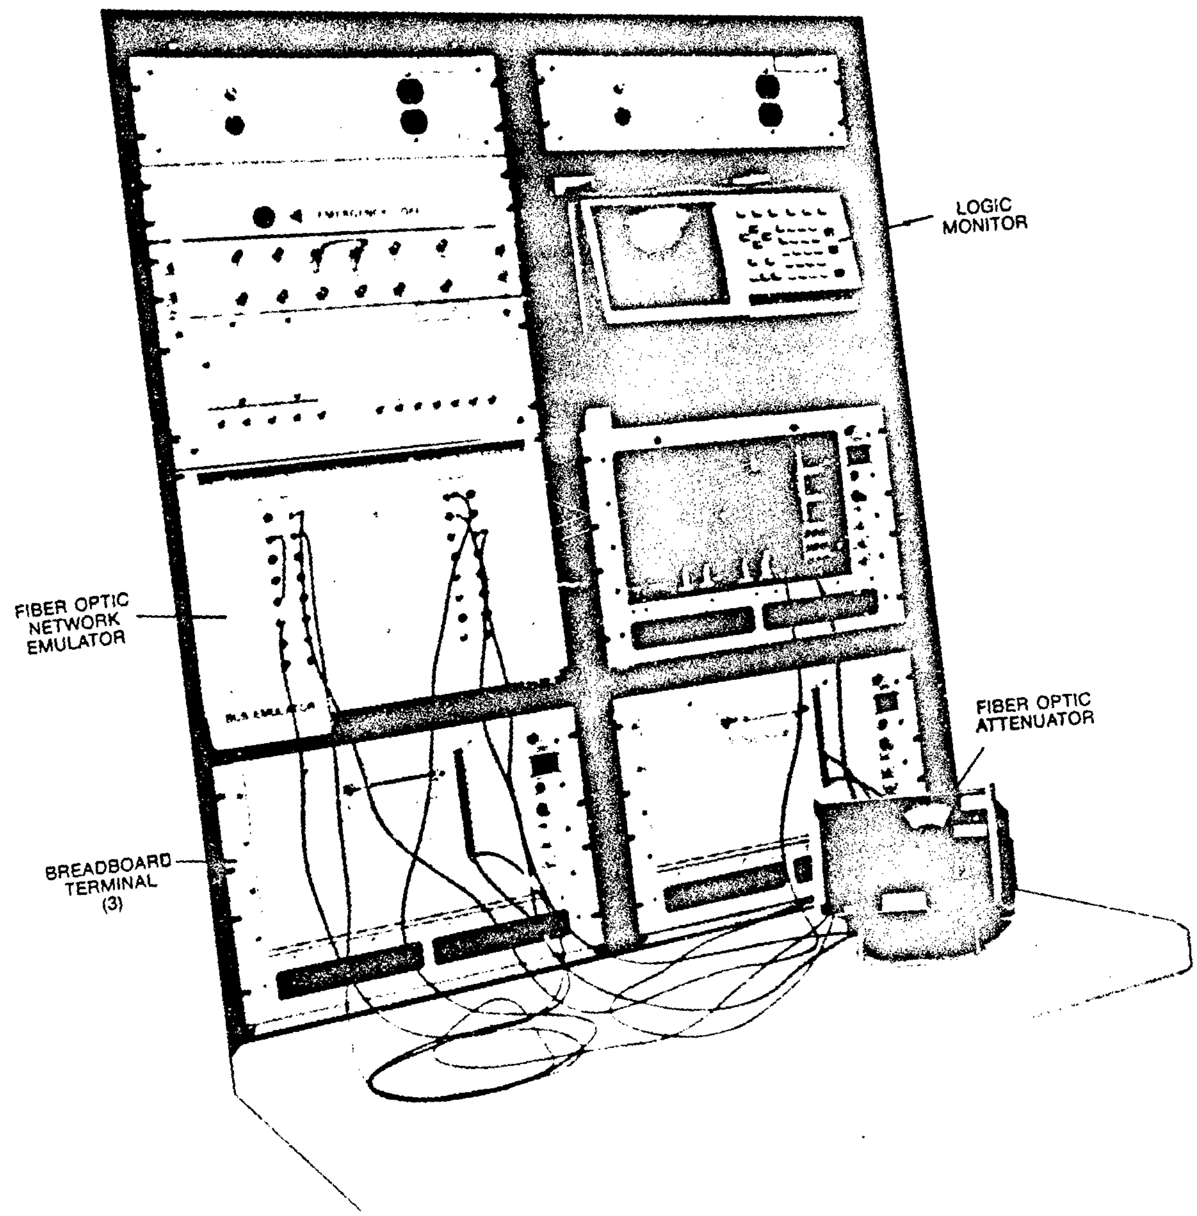

Figure 62. HSDB System Demonstration Equipment Configured with Fiber Optic TRUs and Network Emulator

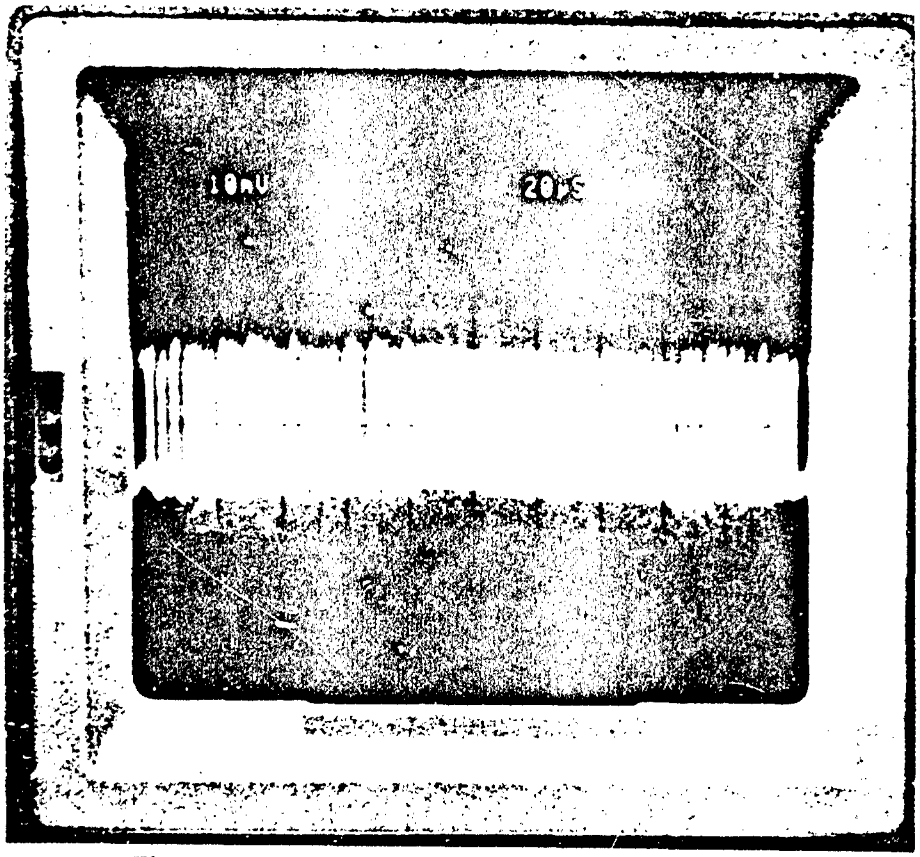

Figure 63. Fiber Optic Token Passing Network in Operation

5.0 DEVELOPMENT OF THE PAVE PILLAR PROTOCOL

Development of the PAVE PILLAR (9) protocol encompassed investigation of the requirements and potential existing solutions, and then synthesis of a protocol which met the needs envisioned for both near-term platforms such as ATF and more technologically advanced applications. This effort was performed under Task IV of the HSB Technology Development contract. After studying the work in this area done by other organizations, it became apparent that there was not one single protocol which is optimum for all of the diverse needs envisioned. For this reason, Rockwell chose as its objective that the protocol:

a. Should meet all requirements envisioned for ATF and the PAVE PILLAR architecture with at least 50% reserve capacity.
b. Should be applicable for a wide range of applications; with between 2 and 64 nodes, intelligent interfaces and unsophisticated interfaces.
c. Could be packaged as one node per SEM-E module by 1989 timeframe.

Based on these broad goals, Rockwell implemented an engineering plan of four phases. The intent was to gradually zero-in on a design which used ongoing protocol work being done by other organizations to the extent practical and also allowed visibility of Rockwell's efforts so that these other organizations might benefit from this program. The intent was to tailor an existing protocol rather than to synthesize an entirely new one.

- **Step 1 was the survey phase.** The intent was to identify needs as envisioned by anyone who might have an application for the PAVE PILLAR HSDB. Out of this phase came the requirements for the HSDB.
- **Step 2 was the baseline protocol design phase.** This consisted principally of identifying candidate protocol alternatives from published and unpublished literature, codifying their characteristics, and selection of a single baseline design.
- **Step 3 was the optimization phase.** Characteristics of the baseline protocol were modified on a sample basis to identify the impact of each potential change on operation of the network. This phase consisted of a series of trade studies, the most significant of which were simulation trade studies.
- **Step 4 was the verification/validation phase.** Prototype HSDB hardware and software were developed to prove-out the simulation results.

(9)PAVE PILLAR HSDB System Specification USAF Contract F33615-83-C-1036, CDRL No. 10-3

From this engineering plan came the PAVE PILLAR HSDB system specification and also the simulation/verification/test data which shows it to be a flexible, reliable, and mature design.

5.1 Requirements Survey

The requirements survey (Step 1) consisted of contacting each ASA contractor and other potential users for as much information as they could (would) supply. In many companies, this consisted of as many as three separate organizations: their ASA program; their ATF program; and their advanced planning organization. From these surveys, a singular requirements document was generated. One of the most significant accomplishments resulting from the survey was development of a consolidated message inventory. This inventory recorded the type of message and the important characteristics listed in Table 16.

Table 16. Message Inventory Characteristics

ITEM	CHARACTERISTIC
a.	Message length – number of 16-bit words
b.	How often does the message repeat?
c.	What is the tolerable delay?
d.	Is the message periodic?
e.	What is the tolerance on periodicity?
f.	Is the real-time at which the message was generated important?
g.	Is immediate acknowledgment required?
h.	Is there more than one destination?
i.	Is the message classified?

At the conclusion of the survey, the message inventory contained more than 500 messages, supplied by six different sources.

Figure 64a and 64b summarizes some of the important characteristics of the message inventory. Note that message length varies from 2 words to 2100 words, with an average of 30 words, and repetition rates are from asynchronous to 60 times per second. The average data rate is 3 Mbps.

Several understandings resulted from this survey which were widely influential during the following phases of the program.

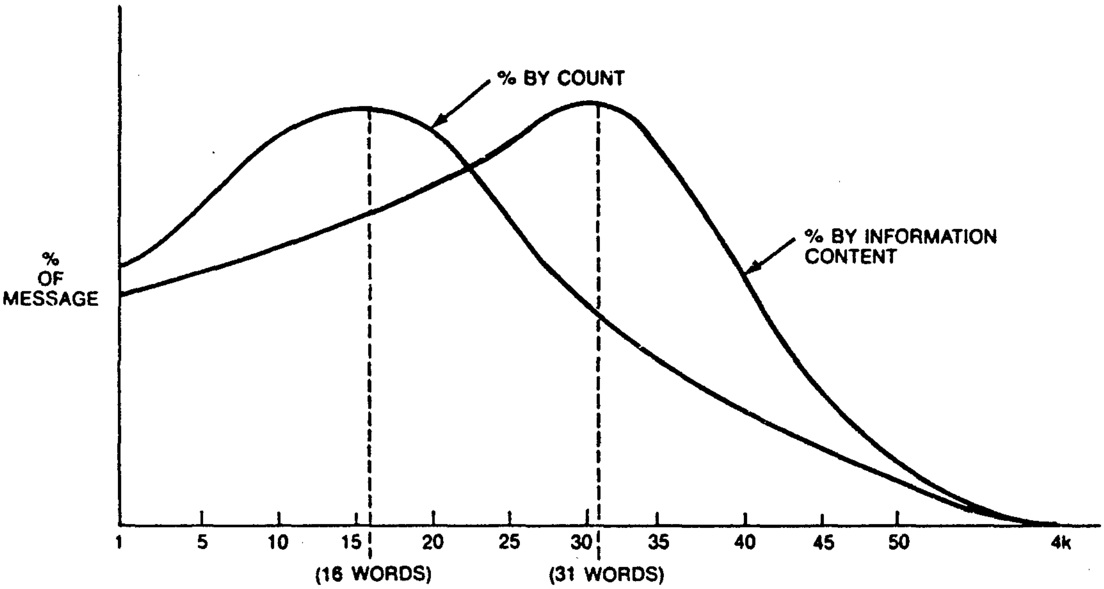

Figure 64a. Message Length Distribution

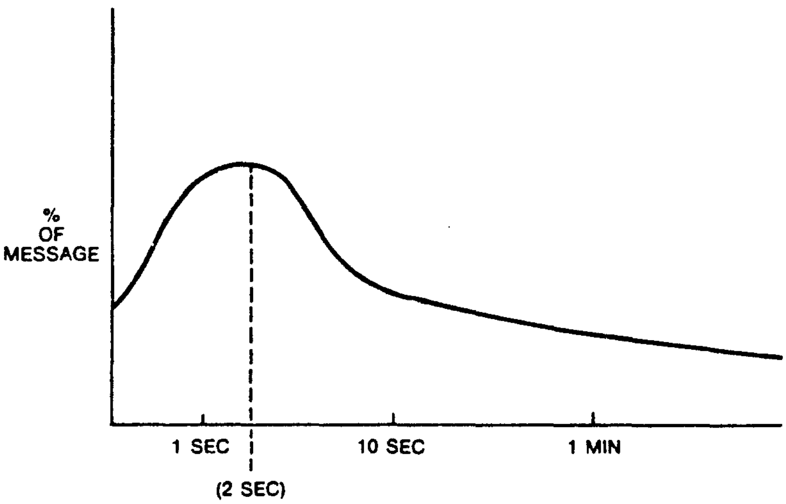

Figure 64b. Message Repetition Rate Distribution

a. 'Modern' aircraft architectures are widely variant. Therefore, Rockwell decided that it would be inappropriate to base design decisions on a narrow set of requirements. Instead, we determined that it was important to provide the systems designer with a variety of options and a method of allowing him to intelligently trade them against one another to arrive at a design appropriate for each specific applications.

b. Aircraft systems designers are overwhelmingly interested in worst-case latency rather than average latency. This rendered useless a large percentage of prior analysis, which was done analytically, and drove us to rely on simulation results instead.

c. Data bus traffic patterns would likely be 'bursty.' Large windows of time could occur with little or no traffic being sent followed by a period where peak traffic might be an order of magnitude greater than the average. The protocol must gracefully handle these peak periods by providing guaranteed latency for high priority messages while maintaining the backlog of lower priority messages.

d. HSDB terminals should contain embedded intelligence to route messages independent of interaction with the local user. To send a message should require only a single simple transaction initiated by the user, to receive a message should require only a single simple transaction initiated by the HSDB terminal.

e. The HSDB terminal should provide a variety of service functions (functions not directly involved with transfer of data across the HSDB). A global reference clock was ranked as first in importance. Distributed monitoring and adaptive tailoring were rated as next most important. The ability to mix secure and clear text data was determined to be unimportant.

Table 17 summarizes the requirements defined for the PAVE PILLAR HSDB protocol. These requirements came from Task I/II results, direction from the Air Force program office, and from the requirements survey. These requirements drove later development effort on the protocol.

5.2 Baseline Protocol Design

When this program began, the industry was blessed with a very rich environment of proposed protocols. Dozens existed on paper, almost none had been demonstrated. Early during the program, Rockwell began to collect literature describing these different protocols from both published and unpublished sources. This provided the basis for later protocol development work. Numerous innovative ideas had been proposed by various members of the LAN community, not all of which were compatible. One of the difficult problems we faced was that an objective analysis of many of the proposed solutions showed little difference in network performance. This meant that many decisions could not be based purely on technical superiority. Those decisions were the most difficult to deal with especially when there were vocal proponents for both sides of the issue. After studying the work done by other organizations and individuals,

it became apparent that there was not one single protocol which directly met all of the program requirements. The direction for our protocol design effort can be stated simply:

1. Select a well developed baseline protocol from among the available candidates.
2. Optimize the baseline for the specific PAVE PILLAR requirements.

Table 17. PAVE PILLAR HSDB Protocol Summary

REQUIREMENT	SOURCE
Data Rate = 50 Mbps	USAF Direction
Number of nodes = 64	USAF Direction
Maximum Physical Separation = 100 meters	USAF Direction
1 to 4096 Word Message Length	USAF Direction
Latency < 10 mS (priority messages)	Survey
Reliability "Beyond Measure"	Survey
Recovery from Network Crash within 1 mS	Survey
Distributed Management	USAF Direction
Broadcast Topology	Task I/II
Coaxial and Fiber Optic Compatible	USAF Direction
Intelligent Terminals	Survey
Token Passing Management Mechanism	Task I/Task II

From more than 30 candidate protocols investigated in some depth features of two were selected to form the baseline PAVE PILLAR protocol, SAE AE-9B/L Draft C and IEEE 802.4. There had been a significant amount of prior work directed toward avionics applications by the SAE AE-9B/L subcommittee. IEEE 802.4 was well documented and tested though the intended application was quite different.

It is important that a protocol be selected that allows as much data throughput as possible, but is is just as important that service be offered for certain messages with minimum

delay. The effectiveness, Q, of the protocol, then, is some function of efficiency, E, and latency, S, so that:

	Q	=	f(E/S)
where	E	=	data throughput bus rate
and	S	=	service interval, (The time it takes for a terminal to once again gain control of the bus after having relinquished control.)

Both the efficiency, E, and the latency, S, are heavily dependent on the message characteristics, as well as the protocol implementation. The average latency can be easily determined for any particular scenario such as those proposed by the SAE. It should be remembered, however, that the above computed latency is an average latency, and that maximum latency can only be computed by defining a worst case combination of busy terminals and/or long messages during a maximum cycle time.

5.3 Optimization of the Protocol

Characteristics of the baseline protocol were modified on a sample basis to identify the impact of each potential change on the operation of the network. This trade study was principally performed using computer simulations of expected HSDB activity.

Selection of a token passing protocol was not without its opponents. Some systems engineers felt more comfortable with a more predictable protocol than token passing allows. They felt that it was necessary that at precisely (or nearly so) a given time, information would be made available. Token passing, by its very nature is not precisely predictable. Information offered from many sources with unsynchronized clocks will at some time arrive in bunches causing temporary network workload peaks. The most extensive trade study performed focused on determining how close to deterministic operation the token passing protocol could be made to perform.

To begin we evaluated the basic SAE AE-9B/L Linear Implementation Task group protocol with no latency control. We also evaluated the impact of overhead size, and of average and maximum message length. Under latency control we evaluated the impact of the token rotation timer settings, of the overuse of priorities, and of the effectiveness of the four level priority system. We also evaluated additional delays caused by such normal but unscheduled activities as adding or deleting terminals. Our final analysis included the evaluation of additional delays resulting from such abnormal activities as error or fault recovery. As a result of this trade

study, we concluded that a token passing protoco: could provide performance closely matching a fully deterministic protocol. This provided technical justification to proceed with the program.

In order to characterize the performance of the network and to examine the wide range of alternatives for all the potential scenarios, it was decided that a protocol and message format simulation tool should be developed. This tool allowed characteristics of the protocol to be varied, and provided performance predictions for:

a. Data throughput rate
b. Latency time for both priority and nonpriority messages
c. Average and worst case traffic distribution
d. Effect of various error/fault recovery methods on latency

Data provided by the simulation tool for various protocol alternatives, in conjunction with other analyses, resulted in the final PAVE PILLAR protocol design.

The success of a simulation approach in performing any trade study is determined by the data base. If the data base provides a realistic scenario then the simulation results will closely predict the behavior of a 'real' system. If the data base is poor then conclusions reached are suspect. This is why much effort went into defining the message data base. Message information came principally from ASA contractors. It was organized in a manner which allowed it to be used in a variety of ways to investigate various protocol parameters. One of the unanticipated characteristics of this data base is that while most messages are presented to the network on a cyclical basis, the instantaneous offered traffic was found to be highly irregular. This characteristic is attributable to our assumption that data sources are not time synchronized. This results in occasional instances where many messages are offered to the network at essentially the same time. Figures 65 and 66 illustrate this. Figure 67 shows a typical scenario. The simulation shows that the average offered traffic rate is 15 Mbps but that the instantaneous (500 μS) rate varies from 0 to 80 Mbps. Obviously, at points where the offered rate is higher than the networks throughput of 50 Mbps, a backlog will be created. This does not mean that messages will be lost, only that messages will be delayed as the backlog is worked off. Figure 66 shows message delay (the time between when the message was offered to the HSDB node for transmission and when it was actually sent) as a function of mission clock time for the same scenario. Note that worst-case message delays are as much as 10 times the average message delay. Rockwell believes that these characteristics closely match the network workload which will be seen on the ATF and other modern aircraft. Characteristics of the data base used for the trade study are documented in paragraph 5.3.1.

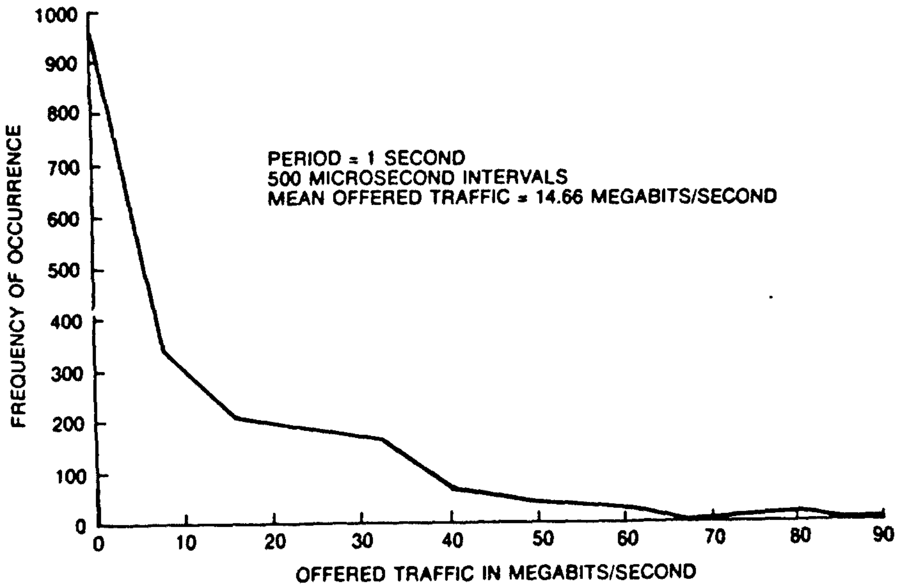

Figure 65. Distribution of Offered Traffic Short Term Rate

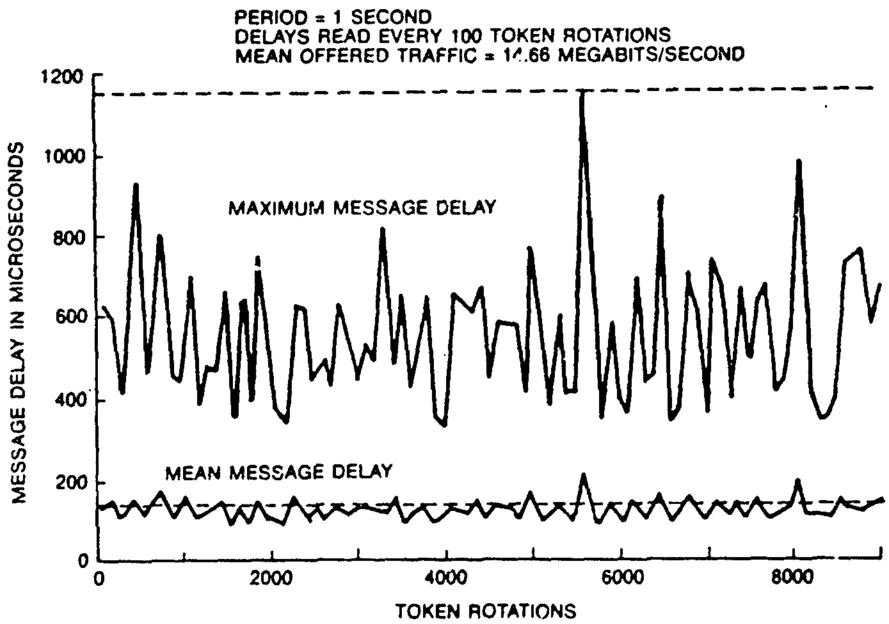

Figure 66. History of Short Term Message Delay

Data Throughput Rate Sensitivity

Data throughput sensitivity refers to the relationship of message delay versus offered traffic rate. Throughput was analyzed in a variety of different scenarios. Figures 67 through 70

show the results of several of these simulations. Each of these figures was generated from the results of several simulations which were identical except for changes to a single variable. The resulting data set was plotted to graphically illustrate the sensitivity of network throughput to that particular variable.

Figure 67 illustrates the throughput sensitivity of the network to offered traffic rate. In this simulation, the standard message data set was used. It includes messages from 2 words to 250 words in length with a mean message length of 30.77 words. This closely models the message data derived from our survey of ASA contractors; the only change being in limiting message length to 250 words. Note that the simulation shows the network to be well behaved (no radical departures from a linear relationship) to above 30 Mbps offered traffic rate. It offers less than 1 mS average message delay for offered traffic rates to 40 Mbps. Worst case message delay is less than 5 mS for offered traffic rates to 40 Mbps. This shows the HSDB is a high performance network, even without the use of a latency control mechanism. It should be noted that this very basic mode of operation meets all performance requirements identified by the ASA contractors during our survey.

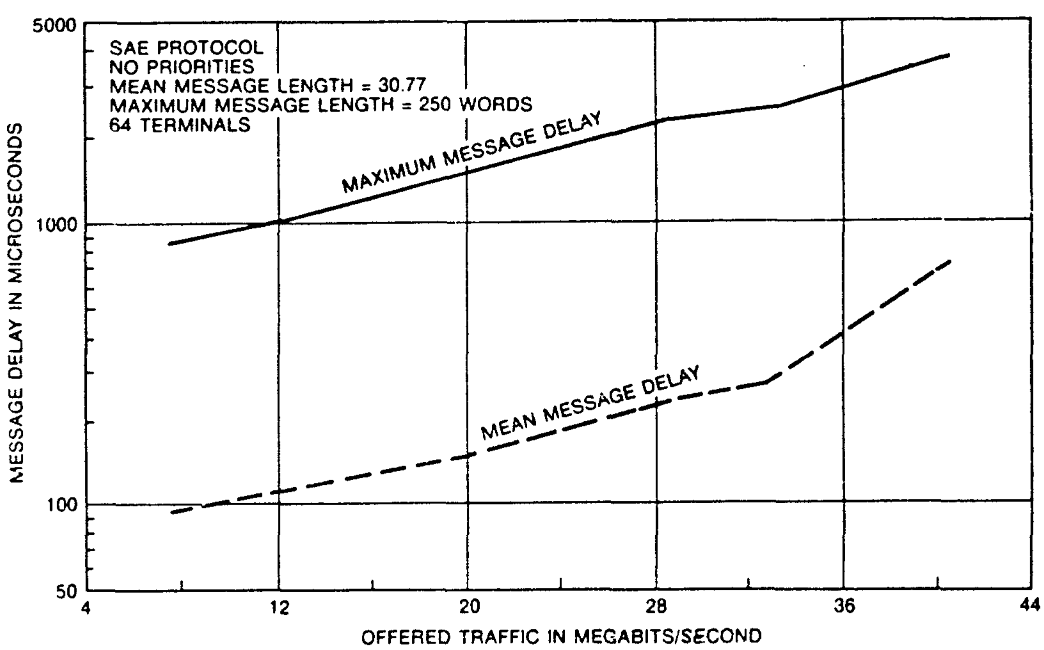

Figure 67. Message Delay vs. Offered Traffic Rate

Figure 68 illustrates the throughput sensitivity of the network to message overhead time. Message overhead time includes terminal propagation time, preamble time, and packet header time. The same data set used for Figure 68 was used except for modification of message

overhead characteristics. Note that performance is relatively unaffected by overhead time. The shortest overhead is SAE AE-9B/L, the longest is that of the contract Task III demonstration protocol. The PAVE PILLAR protocol has a slightly greater overhead than the SAE AE-9B/L protocol because of features added to increase reliability but the performance varies by only 1.25 percent.

Figure 69 illustrates the throughput sensitivity of the network to mean message length. The standard data set was modified to vary the value of mean message length while keeping offered traffic rate and maximum message length constant. The simulation was performed for two different conditions, at approximately the 20 Mbps throughput (the program goal) and at approximately 30 Mbps (50% above program goal). This simulation again shows the protocol to be well behaved. Both mean and worst case message delay increased as mean message length increased but there was no point where performance deteriorated at a rapid pace.

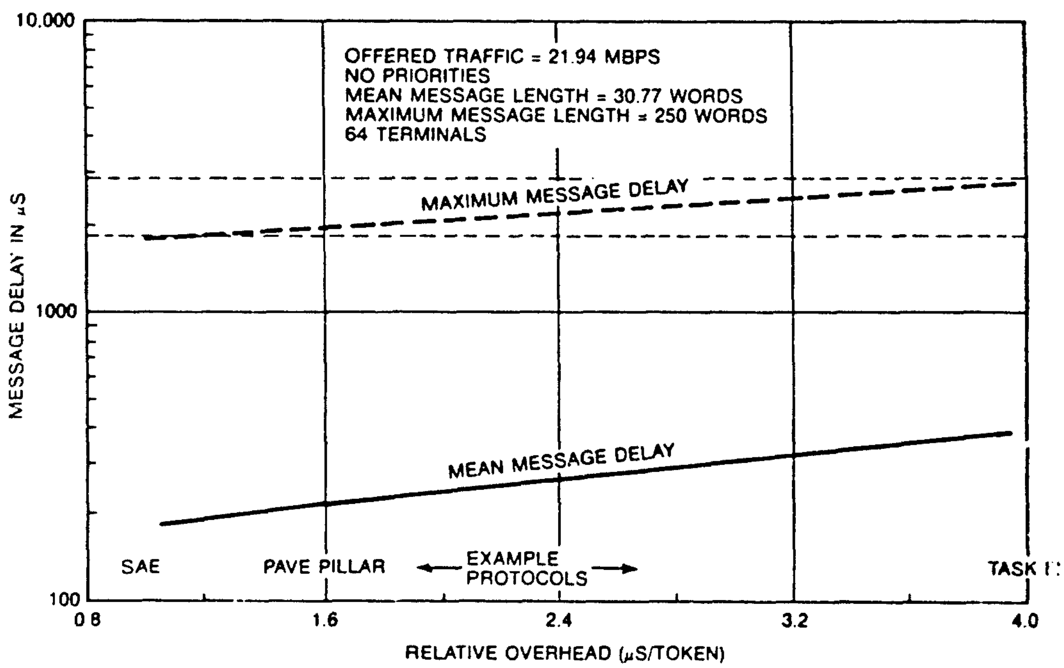

Figure 68. Message Delay Characteristics of Several Protocol Overheads

Figure 70 illustrates the throughput sensitivity of the network to maximum message length. The standard data set was modified to vary the length of the longest messages from 250 words to 4000 words while keeping mean message length at 31 words and offered traffic rate at approximately 20 Mbps. The simulation shows the expected impact of allowing longer messages. Both average message delay and worst case message delay are approximately linearly related to

message length. Even with 4000 word messages allowed, however, the worst case message delay was less than 10 mS, which meets the design goal based on requirements received from the ASA contractors.

As a result of these simulations, we felt comfortable with the base token passing protocol. Message delays were within design goals even without resorting to a latency control mechanism. The protocol seemed relatively insensitive to protocol characteristics. This gave us confidence that the ultimate protocol would work well even though the simulation data base used for optimization was not exactly what might be encountered in real applications. In short, the protocol appeared to be efficient over a broad operating envelope.

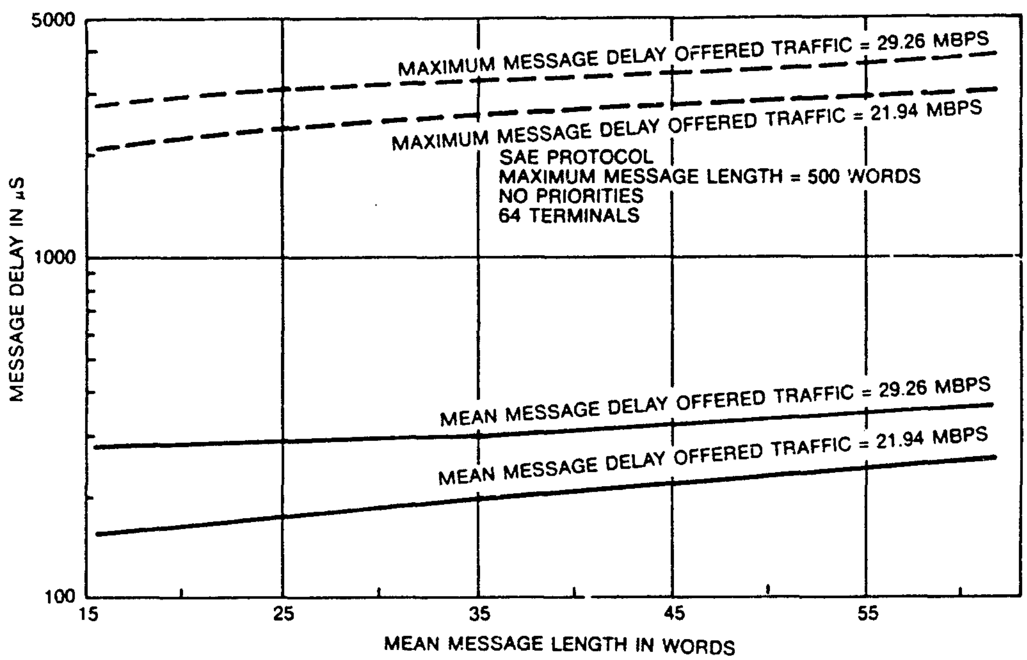

Figure 69. Message Delay vs. Mean Message Length

Latency Control

By any conventional measure, the PAVE PILLAR HSDB is a high performance network even with no latency control mechanism. Latency control was, however, of overwhelming importance to the ASA contractors and was used as a principal criteria for optimization of the protocol. Figure 66 shows the mean and maximum message delay of a typical mission with no latency control mechanism in use. As can be seen, occasionally a high instantaneous offered traffic period will occur resulting in a much above average message delay for a short period of time. The PAVE PILLAR protocol was designed to offer the systems engineer a method of

providing guaranteed low latency for certain critical messages even during these peak traffic periods. This is accomplished by deferring the transmission of other lower priority messages until the network's workload is lower. Delaying low priority traffic is not a problem. That is why it is done--to give preferential treatment to certain critical messages at the expense of those messages for which latency is not so critical.

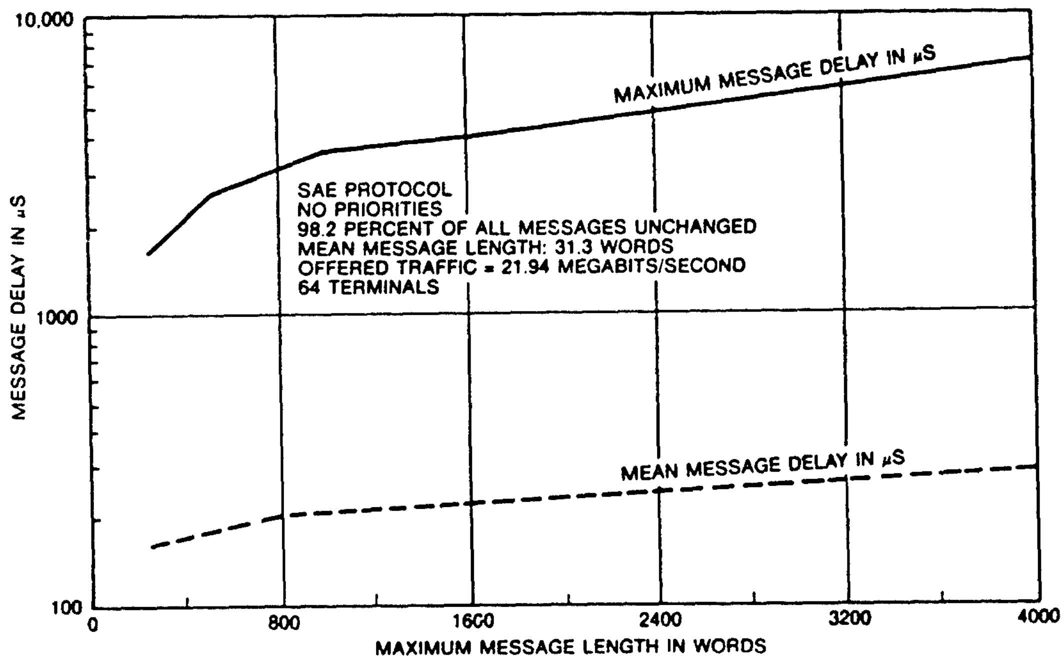

Figure 70. Message Delay vs. Maximum Message Length

Any latency control approach also carries with it a penalty. As network workload increases, more messages are deferred. This creates a backlog which must be worked off. Obviously the use of a latency control mechanism automatically results in lower average network efficiency. Any time all messages are not allowed to go out on the first available token, average bus performance will be degraded. The objective, therefore, of this trade study was to determine (a) what improvement can be expected for high priority messages, (b) what degradation occurs to non-priority messages, and (c) over what operating envelope does the mechanism work well.

The latency control mechanism chosen for use by the PAVE PILLAR protocol uses a set of three token rotation timers (TRTs) similar to those described in the SAE AE-9B/L and IEEE 802.4 standards. This implements a 4-level priority scheme as follows:

a. Priority 0 (P0) messages are always sent (unless maximum packet size would be violated).
b. Priority 1 (P1) messages are sent after P0 messages unless the P1 timer (TRT-P1) was expired when the token was received.
c. Priority 2 (P2) messages are sent after P1 messages unless the P2 timer (TRT-P2) was expired when the token was received.
d. Priority 3 (P3) messages are sent after P1 messages unless the P3 timer (TRT-P3) was expired when the token was received.

All TRTs are tested and then reset each time the token is received. A measure of network workload is provided by noting which timers have not expired when the token is next received. These TRT may be used for latency control.

If the latency required of any specific message is less than the peak values shown in Figure 66, then steps must be taken to prevent those peaks from exceeding the required latency. If one mS maximum latency is necessary, for example, Figure 66 shows that the token rotation timer will not often come in to play. If 500 µS maximum delay is required, the token rotation timer will come into play much more often perhaps as many as half of those 100 token rotation intervals.

The maximum and mean values shown in Figure 66 were derived during a simulation run at about 15 Mbps. Figure 67 shows how these values vary as the offered traffic is increased from about 7 to 40 Mbps. At 30 Mbps, the maximum delays may be expected to reach approximately 2.2 mS . If one mS delay maximum is required, as in the example discussed above, then token rotation timers may be used to assure timely delivery of important messages.

Figure 71 shows that the PAVE PILLAR scheme does exactly what is intended in a two-level priority example. As the token rotation timer setting is reduced, the maximum delays for priority messages are reduced. At the same time, the maximum delays for the non-priority messages are increased. Note that the worst case message delay is approximately equal to the TRT setting until the TRT reaches a quite low value. This makes it simple for a systems engineer to decide what the value of the TRT should be, just set it to the worst case latency required of the priority messages. Operation with 3 or 4 levels of priority is similar. Draft C of the SAE AE-9B/L standard, from which the PAVE PILLAR protocol was derived, used a fixed 2-to-1 relationship between TRT at adjacent priority levels:

$$(\text{TRT-P1}) = 2 * (\text{TRT-P2}) = 4 * (\text{TRT-P3})$$

The rationale for establishing this ratio for token rotation timers was probably in anticipation that this ratio would provide a similar ratio of the maximum message delays for the three levels

of priority groups. Numerous computer simulations demonstrated that this is not the case and that the latency of the lower priority levels is not easily predicted. Rockwell feels that users of the HSDB need to have the flexibility of setting each TRT without regard to a predetermined ratio in order to achieve the desired system response for each priority group. Rockwell also feels that simulation offers the only adequate method of determining optimum TRT settings in such an environment.

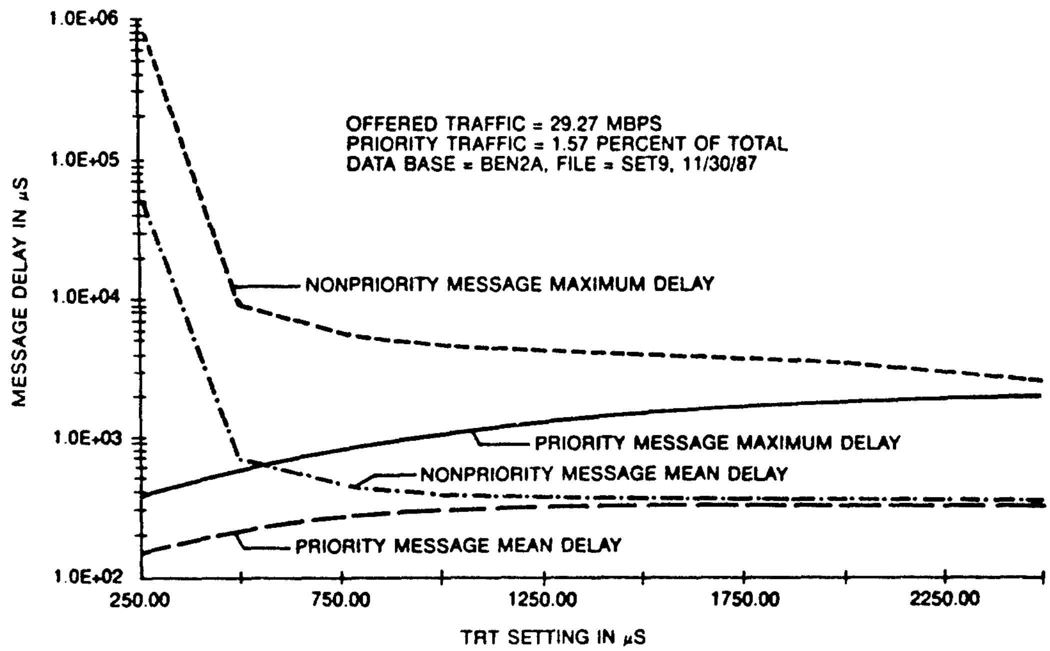

Figure 71. Message Delay vs. Token Rotation Timer Setting

In order to demonstrate the interactive nature of a 4 level priority system, the standard data base was modified and a series of simulation experiments were performed. The total traffic was divided as equally as possible into the four priority groups. The intent of the experiment was to arrive at a TRT setting which would evenly distribute the maximum message delay across all priorities. Table 18 shows the results of this series of simulations.

Simulation #1 shows the performance of a non-priority system for use as a baseline. Simulations #2 through #6 show the performance for different initialization values for each TRT. Simulation #2, for example, uses the default TRT settings defined in Draft C of the SAE AE-9B/L standard. Note that very little difference exists between P0 and P1 performance indicating that the network is not optimized. Other simulations, when viewed as a set, show the interactive nature of the TRTs. None of the simulations resulted in exactly the desired distribution of message delays. Additional simulations with readjustment of the token rotation

timers or the distribution of messages between the priority groups would be needed to achieve the best compromise of the desired results. Our conclusion was that limiting the token rotation timer settings to a ratio of 2:1 severely limits the flexibility of the system design engineer to achieve desired network performance. For that reason, the PAVE PILLAR protocol specifies individually programmable token rotation timers.

Table 18. Iterative Process of Setting TRTs

	Simulation Cycle Number					
CHARACTERISTIC	**1**	**2**	**3**	**4**	**5**	**6**
Mean token rotation time	156	156	156	156	156	156
P0 percent offered traffic (%)	100	24.97	24.97	24.97	24.97	24.97
P0 mean message delay (μS)	261	181	174	179	163	180
P0 max message delay (μS)	1061	562	518	532	512	531
TRT-1 setting		800	470	470	470	470
P1 percent offered traffic (%)		24.88	24.88	24.88	24.88	24.88
P1 mean message delay (μS)		162	156	168	146	164
P1 max message delay (μS)		538	758	785	682	807
TRT-2 setting		400	350	400	300	400
Percent offered traffic (%)		25.16	25.16	25.16	25.16	25.16
Mean message delay (μS)		236	261	222	280	235
Max message delay (μS)		1009	919	970	1074	1044
TRT-3 setting		200	200	170	170	200
P3 percent offered traffic (%)		24.99	24.99	24.99	24.99	24.99
P3 mean message delay (μS)		852	879	943	1013	857
P3 max message delay (μS)		3038	3032	3137	3149	2894

Figure 72 illustrates the impact on network performance of placing an increasing percentage of total network load under latency control. Note that as the percentage of messages given priority increases, the deviation from the latency set point increases in an approximately linear fashion. Note also that this degradation in performance of the priority message class does not result in improved service to the non-priority class. In effect, the network just becomes less and less efficient rather than merely trading performance among message classes. The conclusion here is that latency control should be used carefully. Only messages truly needing guaranteed delivery should be assigned a priority.

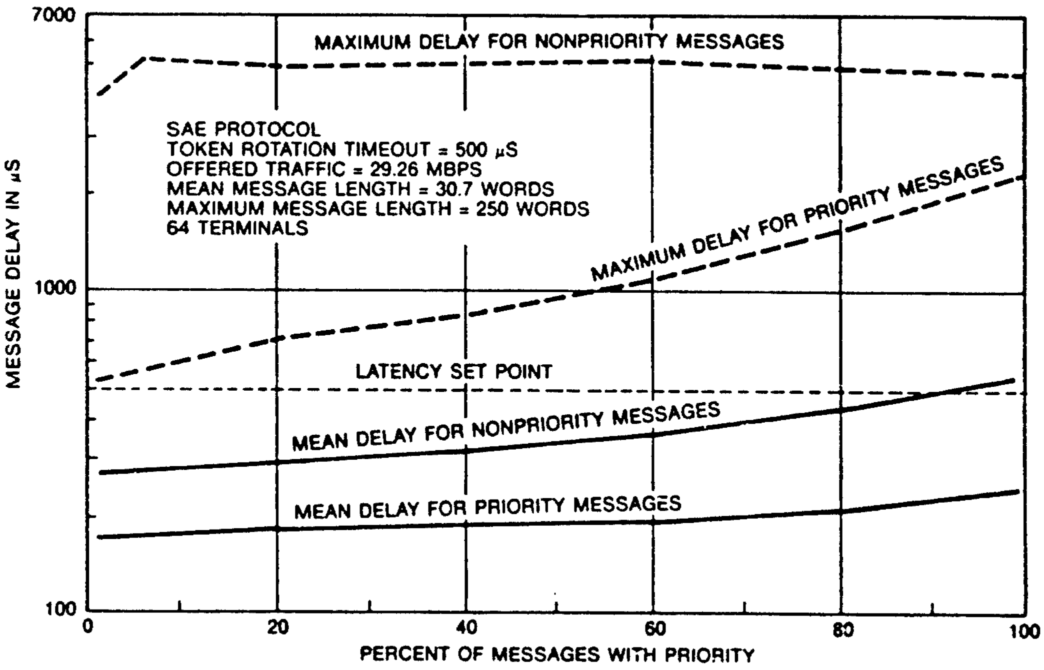

Figure 72. Message Delay vs. Percentage Priority Traffic

5.3.1 Message Data Base

Characteristics of the first 300 messages in the message inventory are shown below:

- Between 30 terminals
- Message lengths: 2-2100 words
- Rates 1-80 messages per second
- Median message length: 16 words
- Average message length: 30 words
- Average data rate: 3 Mbps

Note the difference between median message length and average message length. If we looked only at the average of the inventory of messages identified we would conclude that the average message length would be approximately 16 words, in agreement with the SAE baselines. Note, however, that while only 0.4 percent of the messages are between 1024 and 2046 words long, they represent approximately 28% of the traffic on the bus. As a result the average message length, as determined by dividing the total traffic by the number of messages, exceeds 30 words.

Figure 73 shows message rate vs. message length. This shows that the shorter messages occur more frequently. This characteristic supported the decision to use message length as an

indicator of required latency for some of simulation studies. Shorter messages were assigned high priorities in order to provide shorter latency characteristics.

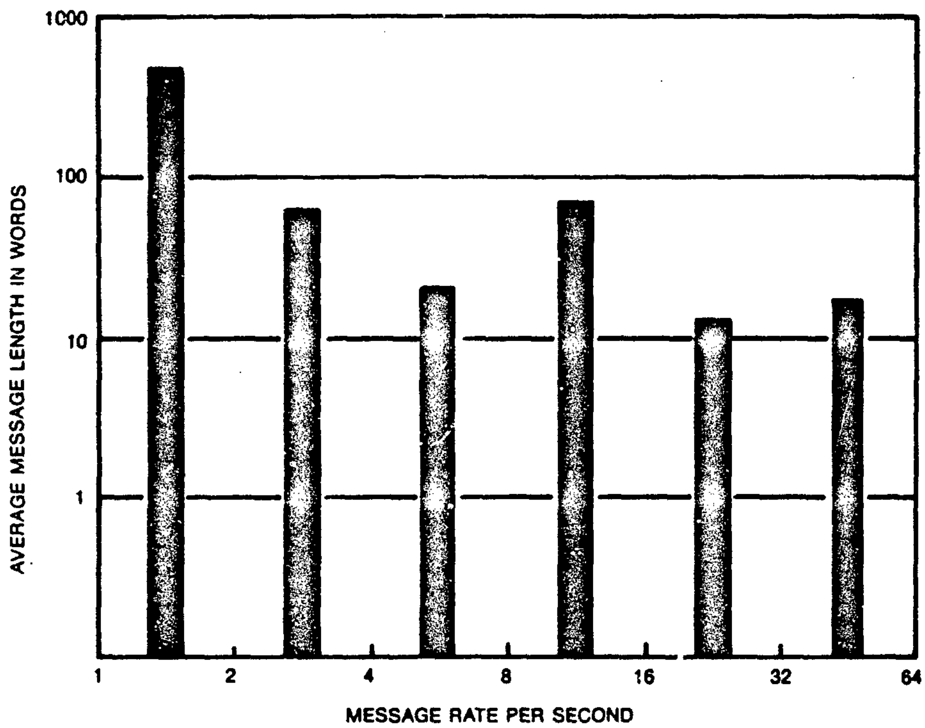

Figure 73. Distribution of Average Message Length vs. Rate

Based on the data compiled during the requirements survey phase, a standard message block was defined which was manipulated in various ways so that a broad range of requirements could be simulated. The standard message block is as follows:

- 157 messages
- 16 terminals
- Rates: 1.25 to 50 messages per second
- Message lengths: 4 to 4000 words
- Average message length: 30 words
- Data rate: 1.8 megabits per second

To simulate a system with 64 terminals, four of the above standard message blocks are used to form a baseline message block of 628 messages and approximately 7.3 Mbps. In the simulation machine, the message rate can be increased by factors of two to six to increase offered traffic data rate from 7.3 to 44 Mbps.

So that we can provide performance information on a wide range of applications, it is important that we are able to vary the characteristics of the traffic load. Among the load characteristics we are able to vary are the following:

- Maximum word length: 250 to 4000 words
- Average word length: 15 to 60 words
- Data rates: 7.3 to 44 Mbps
- Overhead: SAE, PAVE PILLAR, or manually loaded parameters
- Token rotation timers: 250 to 30,000 μS
- Messages allocated to priorities: 1 to 99%

The message data base used for Task IV simulation trade studies consisted of the data sets shown in Table 19.

Table 19. Simulation Message Data Base

FILE	RATE	Avg L (words)	Max L (words)	PO (%)	P1 (%)	P2 (%)	P3 (%)
BEN2A	7.31 Mbps	30.8	250	0	20	20	30
BEN3A	7.31 Mbps	15.4	125	30	20	10	40
BEN5A	7.31 Mbps	61.5	500	0	0	0	100
BEN6A	7.31 Mbps	31.1	500	0	30	20	50
BEN7A	7.31 Mbps	31.2	1000	0	30	20	50
BEN8A	7.31 Mbps	31.3	2000	0	30	20	50
BEN9A	7.31 Mbps	31.4	4000	0	30	20	50
BEN10	7.31 Mbps	15.6	500	30	20	10	40
BEN11	7.31 Mbps	15.6	250	30	20	10	40
BEN12	7.32 Mbps	59.0	250	0	0	0	100
BEN13	14.60 Mbps	30.7	250	0	30	20	50

L - Length

5.3.2 Simulation Tool

Rockwell developed a special purpose network simulation tool as support to the Task IV trade studies. It is an event-driven simulator written in TURBO87 PASCAL and is hosted on an IBM PC/AT. The decision to develop this special purpose tool was made after investigating the wide variety of commercially available simulators and concluding that none could efficiently perform the tasks we envisioned.

Figure 74 is a flow chart which shows the logic of the HSDB simulation tool. To start the model scenario parameters must be set from the console:

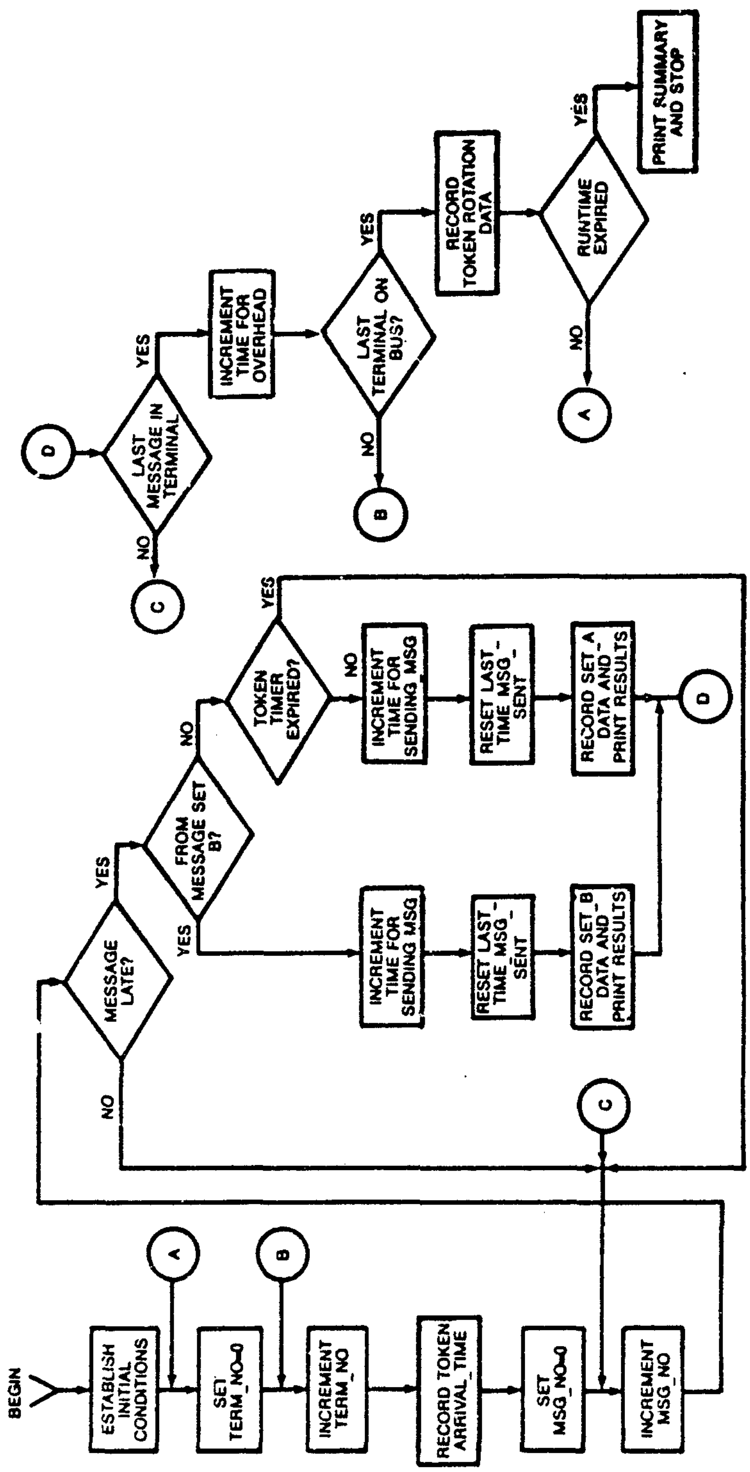

Figure 74. HSDB Network Simulation Model Logic

a. Select data base

b. Select traffic multiplier

c. Select number of terminals

d. Select token timeout

e. Select runtime

f. Load data base

g. Randomize LAST_TIME_MSG_SENT

Default values are provided for each parameter. An option allows the operator to save the scenario for later use. This allows "what if" comparison simulations to be easily run. Once the scenario has been defined, the model is initiated. Beginning with node #0 and message #0, the model checks to see whether that message has been scheduled for transmission since the token last arrived at that node. If not, then message #1, #2, etc. are each checked in sequence. Whenever a message due to be sent is identified, the token rotation timer for that priority is tested. If still active, the model's clock is incremented by an amount of time equal to the transmission time required for a message of that size, simulating transmission of the message. The process continues for each message listed in the message data base for that terminal. When all messages have been checked then the model's clock is incremented by an amount of time equal to the overhead associated with one packet. This simulates the token pass, preamble, propagation time, and terminal response time. The process continues in similar fashion for terminal #1, #2, etc. until all terminals have been processed. At the end of each complete token cycle, the model checks to see if the test runtime has expired. If so, then the simulation results are printed and the simulation terminates. Otherwise, a new token cycle is initiated. Figure 75 is an example of the data recorded for a single simulation. Data may be saved as hardcopy or to a file on disk. This allows the results from multiple simulations to be plotted for ease of analysis. Figure 76 is an example of such a plot.

5.4 Description of the PAVE PILLAR Protocol

The PAVE PILLAR protocol began as a minor modification of the SAE AE-9B/L, draft C protocol. Since that time PAVE PILLAR has evolved, as has SAE AE-9B/L, to a point where present versions of each, while similar, are not interoperable. Rockwell refined several characteristics of the baseline protocol in an effort to optimize it for use on high performance military aircraft. Areas of significant change, and the rationale for the difference are shown in Table 20.

Most differences are attributable to the tightly managed token strategy adopted for the PAVE PILLAR HSDB. These changes were made in order to meet the reliability criteria

```
SIMI4C2
Copyright 1986 Rockwell International Corporation
DATA IDENTIFICATION IS test f
test to compare random data base vs Ben2a
Time = 14:16
Date = 9/21/1987 (Monday)
THIS PROGRAM SENDS ALL MESSAGES FROM A TERMINAL THAT ARE SCHEDULED AT THE
TIME THE TOKEN ARRIVES AND ARE ALLOWED BY TOKEN ROTATION TIMER SETTINGS AND
THE TOKEN HOLDING TIMER. NEW SCHEDULE FOR SENDING MESSAGE ARE INCREMENTED
FROM PREVIOUS SCHEDULE. THE OPERATOR PROVIDES THE OVERHEAD
DATA WHICH INCLUDES GUARD TIME AND TRANSMISSION DELAY
run time in seconds = 1.0
initial random time = 200000
TRAFFIC MULTIPLIER = 1.2
o_hd0 = 1.52 o_hd1 = 2.00 o_hd2 = 1.44
TRT_1 = 380 TRT_2 = 300 TRT_3 = 200
max token holding time selected is 1600
the maximum number of terminals is 64
file:=data_bas
TOTAL TRAFFIC = 857097 WORDS/SEC.
TOTAL MESSAGES/SEC = 26286
percent of traffic in priority 0 = 100.00
percent of traffic in priority 1 = 0.00
percent of traffic in priority 2 = 0.00
percent of traffic in priority 3 = 0.00
p0 min period =        0; max period = 2000000; min length = 0; max length = 300
p1 min period =        0; max period =   95000; min length = 0; max length =  52
p2 min period =        0; max period =   95000; min length = 0; max length =  79
TOTAL NO OF TERMINAL = 64
TOTAL TRFC OFFERED/SEC = 8.57097309249230E+005 WORDS
TOTAL OFFERED MESSAGES = 640
TOTAL OFFERED MESSAGES/SEC = 2.62863348414545E+004
MEAN LENGTH OF MESSAGES/SEC = 32.61
TOTAL OFFERED DATA/SEC = 13.71 MEGABITS
total_len_sec_exp2 = 6.31252478873730E+007
MESS LENGTH STD DEV = 36.5826
NEW MAX DELAYS PRINTED AS FOLLOWS: DELAY, TIME, MESS LENGTH, PERIOD,
TERMINAL #, MESSAGE #
```

Figure 75. Simulator Output Example

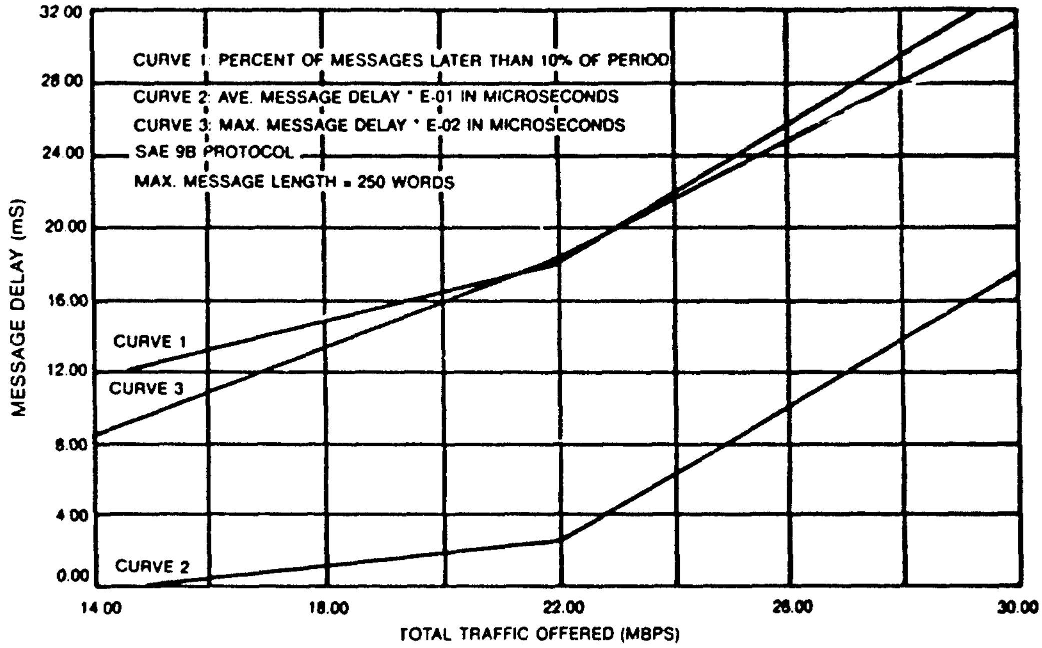

Figure 76. Percentage of Message Delay vs. Offered Traffic Rate

defined in our SOW and those expected by potential users of the network. Other areas of difference came about mostly because the two evolved separately but over approximately the same period of time. Rockwell was driven to make decisions in consonance with the program schedule and was reluctant to change once a decision had been made because of the impact to hardware and software being developed. Rockwell supported the SAE AE-9B/L committee work by regularly describing our findings and decisions at their working group meetings.

Several protocol areas where the SAE AE-9B/L strongly show the influence of work done under this contract are:

a. Wire bus characteristics
b. Programmability of the TRT
c. Redundancy approach
d. Concatenated messages

At the time this final report was written, the SAE standard had not yet been published.

Table 20. PAVE PILLAR vs. SAE AE-9B/L

CHARACTERISTIC	COMMENTARY
Initialization	PAVE PILLAR uses a more tightly controlled initialization approach
Fault Recovery	PAVE PILLAR uses a topology map to speed recovery from error and to allow alternative paths through the network
Token	PAVE PILLAR verifies that a received token is from the expected source terminal
Message Format	Similar but different

5.4.1 Summary Of The PAVE PILLAR HSDB Protocol

The HSDB network consists of a set of nodes connected to a common media as shown in Figure 77. The media may consist of either coaxial cable or optical fiber and interconnects between two and sixty four nodes. Each node is supported by a single HSDB terminal and services one or more users. Information is transferred through the network in the form of Manchester encoded digital data packets. Each packet is transmitted to the network by a single terminal and is received in approximate real time by all terminals (including the transmitting terminal) of the network. A token passing protocol provides synchronization of the network, ensuring that only a single terminal has the right to transmit at any time. Management functions embedded in each terminal provide network initialization, recovery from network faults, and general monitoring of network health.

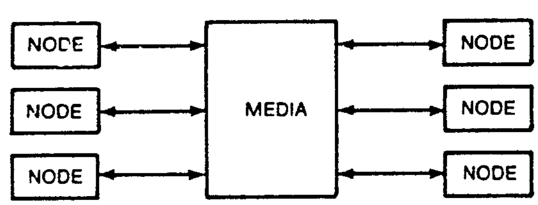

Figure 77. HSDB Network Generalized Topology

Packet Construct

Packets transmitted on the network consist of a sequence of Manchester signaling symbols which constitute one or more messages plus control fields. Figure 78 illustrates the construct requirements of a valid packet.

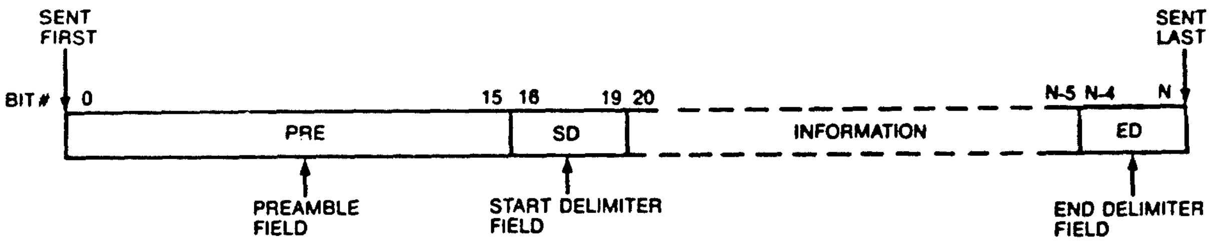

Figure 78. Packet Construct Requirements

a. Preamble (PRE) consists of 16 bits of Manchester Logic "1" symbols at the beginning of the packet.

b. Start Delimiter (SD) consists of a unique 4-bit invalid Manchester symbol occurring immediately following the last bit of the PRE field.

c. End Delimiter (ED) consists of a unique 4-bit invalid Manchester symbol occurring as the last 4 bits of the packet.

The information field contains one or more messages of the form described below. If more than one message is included in the packet, each is separated by a concatenation delimeter field which is identical to SD immediately followed by ED. A maximum packet size of 68920 bits is allowed which accommodates a maximum of 40% data words.

Data Message Construct

Data messages included in HSDB packets are constructed from a sequence of fields described below, and illustrated in Figure 79.

a. Frame Control (FC) consists of 8 bits which define the message type. Unique message types are defined for token messages, data messages, time synchronization messages, and maintenance messages.

b. Source Address (SA) consists of 8 bits which contain the physical address of the terminal which sent the message.

c. Destination Address (DA) contains 16 bits which define the network address to which the message is being sent. Three addressing modes are allowed: (1) terminal physical address, (2) subaddress, logical (multicast) address, and (3) broadcast.

d. Word Count (WC) containing 16 bits which equate to the binary number of 16-bit words which are contained in the data field(s) of the message.

e. Data consists of a quantity of 16-bit words which contain the information to be transferred by the message. The data field of a message with more than 255 words is split into two or more subfields of 255 words, each of which has a frame check sequence field appended.

f. Frame Check Sequence (FCS) contains a 16-bit cyclic redundancy check (CRC) word computed against the previous fields of the message. The CCITT-CRC-16 polynomial is used to calculate the FCS.

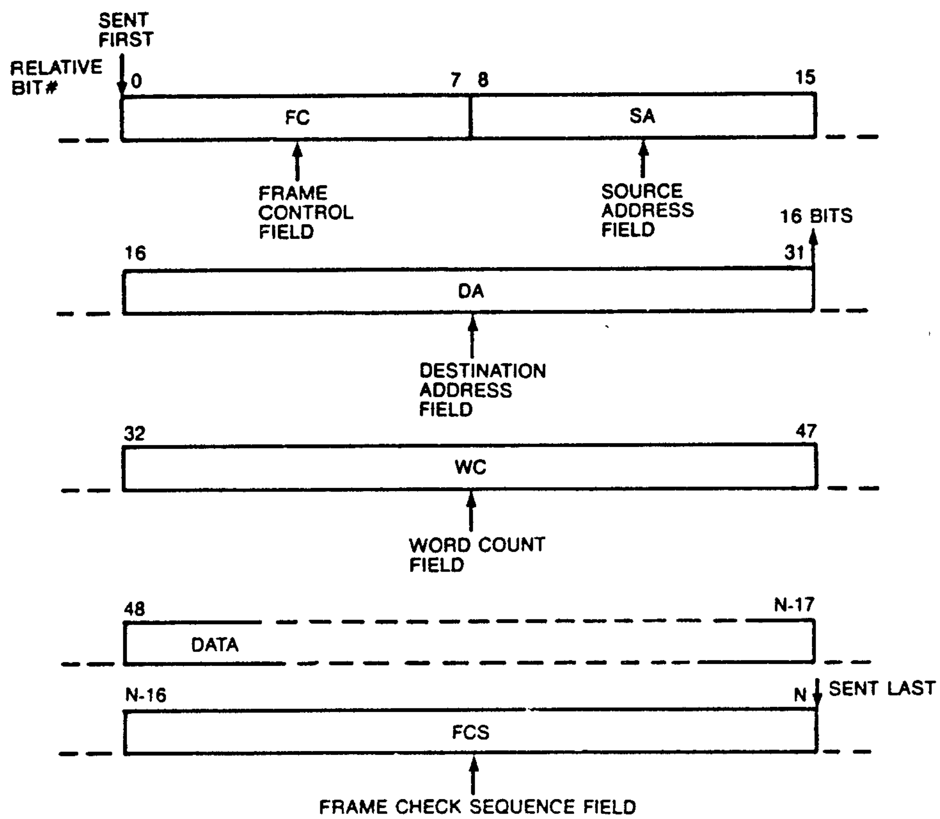

Figure 79. Data Message Construct Requirements

Token Message Construct

Token messages are identical to data messages in construct except for the following fields:

a. The DA field contains only the 8 bit physical address of the terminal to which the token is being passed.

b. The WC field is not present.

c. The data field is not present.

Figure 80 illustrates the construct of a token message. Unique token messages are defined for a normal token, exit token, claim token, solicit entry token, request entry token, and pass Ringmaster token message.

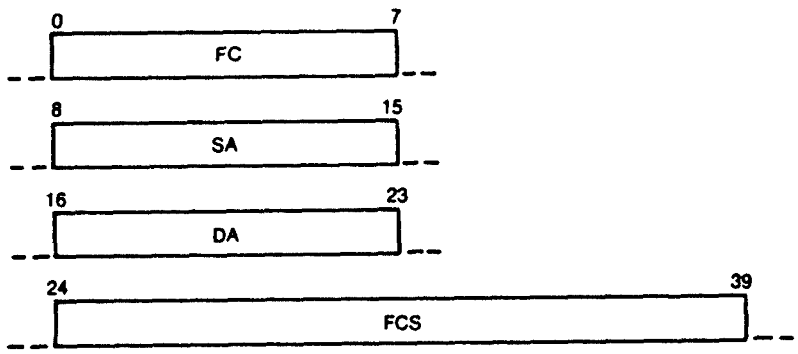

Figure 80. Token Message Construct Requirements

Normal Network Operation

The PAVE PILLAR HSDB protocol implements a token passing protocol to govern access to the network and to coordinate data flow. There are those who believe that a designated HSDB node, for example a centralized controller, should be in charge of initializing and controlling the network. As a result of early protocol design activities, Rockwell felt that token passing protocol would provide the required latency. This obviated the need for centralized control mechanisms. Centralized control was defined for network initialization, however, since it provided improved network reliability and recovery characteristics.

Normal network operation consists of each terminal sending a token message to a specific successor terminal when it has finished transmitting messages held in its queue, if any. A typical sequence of normal network operation is illustrated by Figure 81. The functions required in each terminal include:

a. The ability to receive, decode, and validate (recover) a token message from the network.

b. The ability to recognize that the token source address matches the address designated as its predecessor and that the token destination address matches its own physical address.

c. The ability to transmit a properly encoded and formatted token message to the network upon receipt of a valid token.

The affect of the token passing mode of operation is to create a logical ring control architecture superimposed on the broadcast format physical architecture of the network. Access

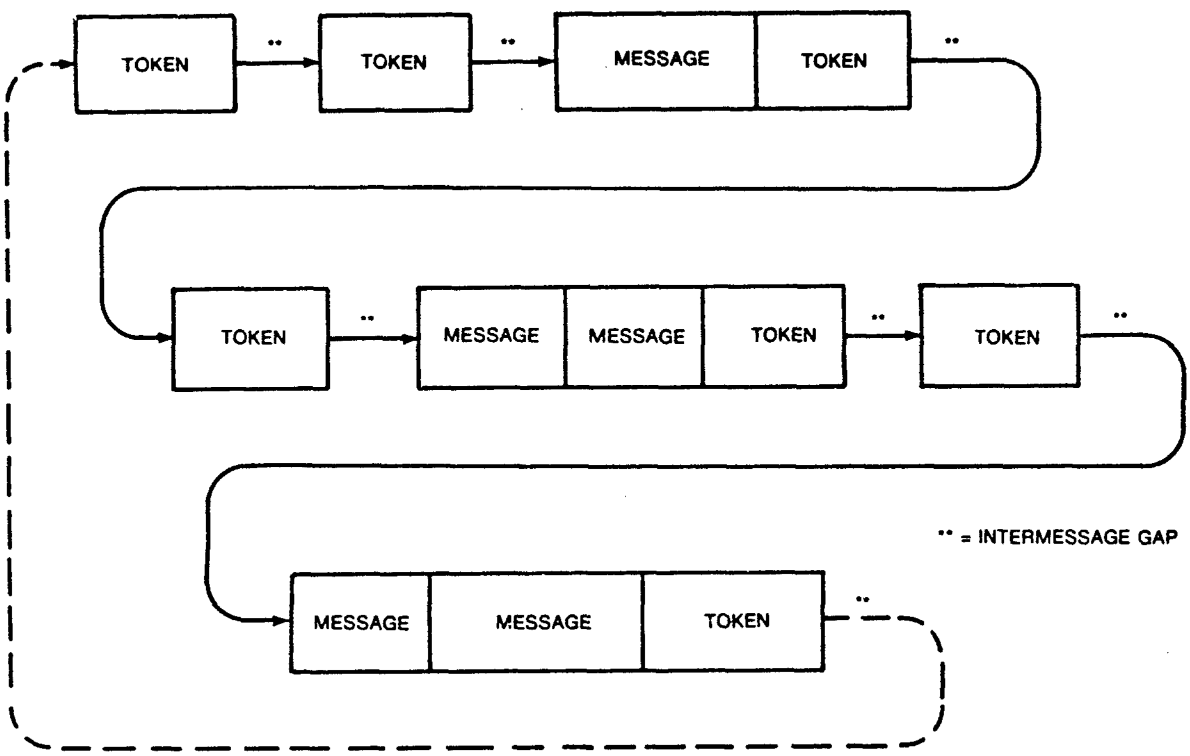

Figure 81. Token Passing Information Transfer Format

to the network, i.e., the right to send a message, is granted sequentially to each terminal by its successor in the logical ring.

Initializing the Network

Initialization refers to the process of defining the logical ring whenever power is applied to the network. Rockwell developed the PAVE PILLAR protocol around the distributed approach proposed by SAE but modified it to provide a mechanism wherein a single terminal, designated the RINGMASTER, establishes the logical ring. This approach supported the tightly managed token management mechanism. A simplified description of the initialization process is provided below. The detailed requirements are described in the PAVE PILLAR HSDB system specification.

At all times each terminal monitors the network for activity, defined as three signal polarity transitions recognized over 4 bit-times. Whenever inactivity is detected, a network inactivity timer (NIT) is enabled. The NIT decrements from a preset (programmable) value until either it is reset by detection of valid network activity or else it reaches zero. Upon reaching zero the terminal assumes that the network is not in operation and attempts to start it by the following procedure:

a. A specialized CLAIM_TOKEN message is broadcast on the network. Each terminal is assigned a CLAIM_TOKEN message of a unique length, determined by the terminal physical address.

b. Upon completion of the CLAIM_TOKEN message, and following a short delay to accommodate network propagation time and settling time, the terminal listens for network activity. Upon monitoring a quiet network the terminal assumes that it has the highest physical address of all active terminals and declares itself to be the RINGMASTER terminal. All other terminals terminate the claim token process.

c. The RINGMASTER terminal polls every other terminal address using a specialized SOLICIT_ENTRY message. Terminals responding with a REQUEST_ENTRY message are linked into a logical ring, terminals not responding are bypassed.

d. When the polling sequence has been completed the RINGMASTER initiates network operation by passing the token to its successor in the logical ring.

Entering/Exiting An Operating Network

Several different methods for adding and deleting nodes while the network is in operation were studied. The approach chosen for the PAVE PILLAR HSDB protocol is similar to that proposed by the SAE in that each station is responsible for checking, at routine intervals, whether inactive stations desire service. This is done in coordination with two timers, the ring maintenance timer (RMT) and the solicit entry timer (SET). The RMT measures network activity and inhibits the process of adding terminals while network activity is high. The SET determines how often the terminal will allow new nodes to be added.

The process of adding a terminal is a subset of the process of initializing the network from powerup; each terminal polls (sends a SOLICIT_ENTRY message) to one address (if any) between its own physical address and that of its successor. If that terminal responds with a REQUEST_ENTRY message, it is spliced into the logical ring by the soliciting terminal and is next to receive the token.

The process just described differs from that of the SAE AE-9B/L standard in two ways:

1. PAVE PILLAR solicits only a single potential new terminal for each assertion of SET rather than the entire address gap as SAE specifies. This was done to limit network latency uncertainties which could result from an unbounded process.
2. The RMT function is not included in the SAE approach. Rockwell simulations showed a potential existed for adding unacceptable delay to priority messages existed if terminals were added during high traffic periods.

A terminal can exit from an operating network in either of two ways.

1. The station going inactive may so indicate to its predecessor by use of the exist token. The predecessor will then (and thereafter) pass the token to the successors successor terminal effectively bypassing the exiting terminal. Rockwell chose this approach strategy in order to support rapid recovery in the case of planned network changes.
2. If a terminal suddenly goes inactive, its predecessor, which normally passes a token to it, will hear no response. After a second try delayed by a period of time determined by the terminals response window timer (RWT), the predecessor terminal will try to pass the token to the next station in the logical ring and continue to try consecutive stations until a station responds. Rockwell chose this strategy in order to minimize the impact to network operation, whenever one or more terminals suffers a catastrophic failure.

This is a more complex recovery process than that defined by the SAE standard. SAE simply requires that the terminal holding the token hunt for a new successor whenever a token pass fails. Rockwell studies showed that this simplistic approach could cause unacceptable network disruption during periods where a number of terminals went inactive simultaneously, such as might occur during battle damage conditions. Under battle damage conditions it is especially important to have the network operating efficiently so that systems reconfiguration can be accomplished.

Fault Response

Recovery mechanisms must be supplied for all abnormal operating conditions which may adversely affect operation of the network. The NIT function provides a catch-all recovery means for all lost tokens (token holder no longer functions), and is likely a good feature of the bus whether or not it is used for adding or deleting stations from the bus. Rockwell chose a conservative self check approach wherein each terminal is smart enough to determine whether or not it is functional to a high level of confidence. If not 100% certain that it is functional, it must go offline so as to not interfere with network activities. Stations that malfunction in such a way that they are not detected by embedded detection methods must also be dealt with. Some kind of monitor function and a complementary recovery function must be provided. The monitor function for the PAVE PILLAR HSDB, is on the RINGMASTER terminal. The NIT in the RINGMASTER terminal reacts more rapidly than the similar function in all other terminals to allow the RINGMASTER to initiate recovery prior to initiation by any other terminal.

Latency Control

The most direct performance criteria for protocol evaluation was worst case latency. This was supported by the survey; all respondents felt that some form of latency control was needed. The PAVE PILLAR HSDB protocol uses a token rotation timer priority mechanism for latency control. This is not a universal approach. Many protocols, including the SAE AE-9B/L, use priorities as a mechanism for sharing access to the network. This is of concern for networks for which there is not a common system engineer (commercial packet switched networks for example) and where each application is "selfish" (i.e., will use network bandwidth for its purposes to the exclusion of all other applications if allowed to do so). Aircraft networks will, however, have a single systems engineer in authority. This engineer will define, across the entire application set, what each priority is to be used for and which messages are to be sent with each priority. There will be no 'selfish' applications in such a network. The priority mechanism may then be used for latency control, a subject of overwhelming importance to potential HSDB users.

The PAVE PILLAR HSDB protocol defines a 4-level message priority hierarchy implemented using token rotation timers (TRT). During operation each terminal schedules messages using a first-in-first-out strategy at each priority level. P0 (highest priority) messages are always sent, within the maximum packet length constraints established for the network. Messages at lower priorities (P1, P2, P3 respectively) are sent only when allowed as determined by the following process:

a. All token timers are initialized to a predetermined value at network initialization. They decrement from their initialization value until they either reach zero count or are reset.

b. Whenever a terminal receives the token it saves the status (active or zero) of each timer, then re-initializes them all.

c. The terminal sends its P0 messages, if any, then checks the status flag for TRT-P1 (the active/zero state saved when the token was received). If the flag shows active, P1 messages are sent (if any).

d. The TRT-P2 flag and the TRT-P3 flag similarly govern transmission of messages at these lower levels of priority.

e. When either the transmit message queue has been exhausted or the maximum packet length has been reached the terminal sends the token to its successor.

Any priority or latency control scheme which offers improved latency to some messages will do so at the expense of the remaining messages. While latencies may not be as critical for some messages as it is for others, some messages cannot tolerate increased latencies. The PAVE PILLAR protocol accommodates 4 levels of priority. Rockwell recommends, however, that

latency control only be used for very critical applications and that a 2 level priority system is adequate for almost all applications.

Probably no other factor directly affects network latency as does message length. The PAVE PILLAR protocol provides direct capability for 4K word messages. If long messages, like 4K words, can be broken up into four 1K word, or eight 512-word, or sixteen 256-word messages, the result will be progressively reduced latencies for the basic network operation. The overall throughput of the network will not be significantly affected; but the latency for other shorter messages will be significantly improved. It is clear that message length of a system should be limited to the length of the longest message for which latency is critical. Rockwell recommends that messages on critical networks be limited to 256 words or less in order to allow the latency control mechanism to function effectively.

Message Acknowledgment

If acknowledgment is required that a message was received by its intended destination user, that user must generate a response message to the sender. The protocol does not contain an automatic notification function. Suspending network operation to automatically acknowledge message receipt would increase latency in an open ended fashion. Most surveyed systems designers felt that modern decoupled architectures could be made to function reliably without the need for immediately acknowledged messages, and that when acknowledgement was required it could be managed user-to-user within the constraints of normal network operation.

Network Reference Clock

The PAVE PILLAR HSDB protocol supports a network wide reference clock function. The intent is to provide a relatively accurate clock which is available to all network users for time tagging messages, etc. Each terminal has an embedded clock oscillator and clock timer register. Across the network one terminal will be designated the TIMEKEEPER and will issue periodic synchronization messages which broadcast the present value of its timer register. All other terminals, upon receipt of the time synchronization message, update the value of their clock timer register to match that of the TIMEKEEPER. Each terminal contains logic which determines when it should be operating as TIMEKEEPER. This allows the TIMEKEEPER function to be assumed automatically should the TIMEKEEPER fail or go off line.

SAE adopted a similar but different reference clock timer function for their standard following a working group meeting where Rockwell described the prototype PAVE PILLAR HSDB protocol.

Redundant Network Interconnect

A dual-path redundant network is the preferred interconnect for the network although a single-path network is allowed. The purpose of redundancy is to provide assurance that every message on the network will be received without the need to reconfigure the physical or logical interconnect of the network. The impact of this approach on the protocol and on terminal hardware requirements is significant. To accommodate the redundant network paths, terminals both transmit and receive on both network interconnect paths (arbitrarily called "A-path" and "B-path") at all times. The message first arriving at the designation terminal is recovered unless it contains one or more errors, in which case the later arriving message is recovered. This approach requires minimal redundancy of terminal transmitter functions but requires two independent receivers plus additional logic to coordinate the recovery of the message from either channel. This redundancy approach can, however, be accommodated within the packaging goals established for the program (SEM-E).

Maintenance Functions

Each terminal contains functions which measure and report its own health and that of other network terminals. In addition to traditional BIT functions the terminal accumulates statistics data which allows analysis of its own operation in the network and also of the network as a whole.

The requirement to accumulate terminal and network statistics data was quite controversial among potential users represented within ASA and SAE. Many thought that the requirement imposed cost, reliability, and complexity disadvantages which outweighed any benefits to be realized. Others had quite the opposite opinion, that the stated requirement did not meet minimum needs for efficient network maintenance operation. Rockwell chose the set of requirements defined in the specification based on our analysis of data needed to properly maintain an aircraft in the field, assuming 2-level maintenance. We recognize that this function is of questionable value during the course of a mission.

Time Tagging

The high-speed data bus was designed to offer a path between two users that traditionally were directly connected. In traditional systems, information was delivered in real time, with no delays, or at least with delays that were consistent from message to message. The high-speed data bus will deliver information in which the delays may vary from message to message. Where this time is critical it may be necessary to include a time tag with the information. It can be argued that time tagging is necessary because the bus cannot function in real time and that the bus should provide that function. However, all users will not need time tagging and this capability should be included only on an as needed basis. The protocol contains no capability to

automatically time tag messages. Rockwell proposes that time tagging be implemented as a user-to-user function where needed.

5.4.2 Evolution Of The PAVE PILLAR HSDB Protocol

Version 1 of the PAVE PILLAR HSDB system specification was released in January 1986 in conjunction with Task IV SRR. The protocol described in Version 1 could be characterized as a minor superset SAE AE-9B/L, draft C.

Version 2 of the PAVE PILLAR HSDB system specification was released in June 1986, in conjunction with the Task IV PDR. The protocol was essentially an expansion of Version 1 with additional technical detail. In Version 2, the addressing approach for content addressed messages was changed to make it compatible with the method adopted by the SAE AE-9B/L committee. This used a 16 K-bit message filter which could be loaded into a terminal either via the HSDB or from its local user.

Version 3 of the PAVE PILLAR HSDB system specification was released in October 1986. The protocol description was changed little from Version 2. The major change involved the redundancy approach which was changed to the dual-path synchronous method. This brought the design into consonance with that planned for the SAE AE-9B/L standard.

Version 4 of the PAVE PILLAR HSDB system specification was released in February 1987, following Task IV CDR. The two changes of significance from Version 3 resulted from direction received from the Air Force Program Office as a result of the initial JIAWG meeting. These were:

1. The reference clock requirements were relaxed to eliminate the need for a precision clock and correction algorithm.
2. The initialization approach was change to "clear token" approach which had been proposed for SAE AE-9B/L.

Rockwell agreed to these changes with the comment that using the "clear token" initialization approach would eliminate the possibility of using the protocol on distributed coupler topologies such as linear tapped bus and multiple stars. The system specification contains descriptions for both the original collision initialization approach and the clear token approach since the newly adapted method would not work with coaxial implementations of the HSDB. Following publication of Version 4 the clear token approach was abandoned by SAE.

5.5 Verification/Validation of PAVE PILLAR HSDB Performance

The PAVE PILLAR protocol design was validated by fabrication, test, and demonstration of two generations of prototype HSDB networks. The breadboard terminal design implemented the core protocol functions of Version 3 of the HSDB system specification. A three-terminal HSDB network was demonstrated in March 1987. The following protocol functions were shown:

a. Establishment of the logical ring from a cold start
b. Transmission and reception of messages
c. Token passing
d. Recovery from a lost token
e. Adding nodes to an operating network
f. Critical terminal timing

The brassboard terminal was designed to implement the full functionality described in Version 4 of the specification. Because of program funding constraints only a single brassboard terminal was fabricated. This did not allow demonstration of an operating network but a majority of the Version 4 functions were validated using laboratory test equipment.

5.5.1 Breadboard Terminal

The breadboard terminal design was intended to verify core protocol functions and to validate simulation based performance specifications. Figure 82 shows the breadboard terminal and the circuit cards which comprise it. The major subassemblies (cards) are:

TRU - One fiber optic transmitter and one fiber optic receiver
RTXM - Receiver/transmitter high speed logic
APM - Protocol logic
AIC - Input controller logic

The breadboard terminal interfaces with a Sperry 1631 computer system which simulated the HSDB user. In this manner an operating HSDB network, complete with simulated users, was demonstrated. The breadboard terminal was designed to validate the core protocol functions using a non-redundant network configuration. The architecture is essentially the same as described below for the brassboard terminal but was implemented entirely using discrete logic devices.

5.5.2 Brassboard Terminal

The brassboard terminal was designed to implement the full protocol described in Version 4 of the PAVE PILLAR HSDB system specification. Figure 83 shows the brassboard

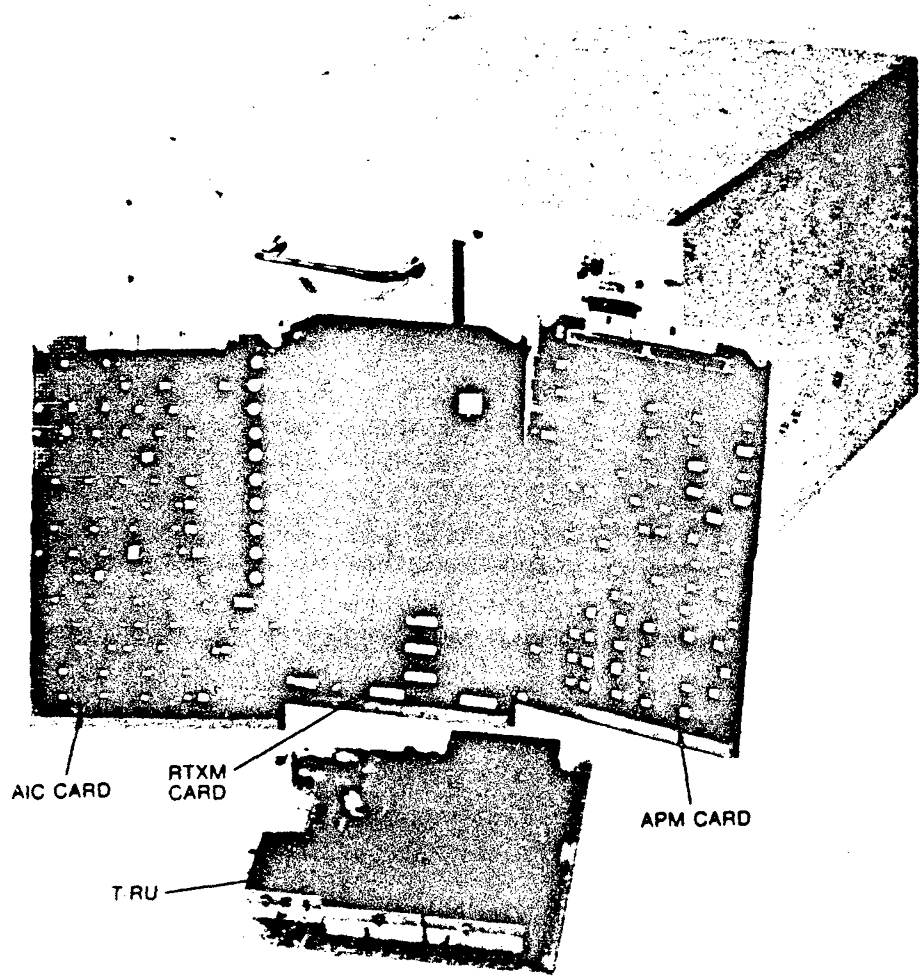

Figure 82. Breadboard Terminal

with cards extended. Two major logic assemblies, the HSDB card and the message handler (MH) card, and two TRU assemblies comprise the functional electronics. Part I (Development) and Part II (Fabrication) specifications for the brassboard terminal have been delivered as data items.

5.5.2.1 Architecture

Figure 84 shows the architecture of the brassboard. HSDB Rx ports provide a path through which packets from the HSDB network may be received. Two ports are provided, one for the A-path and one for the B-path. Both are accessible via a front panel access slot of the terminal.

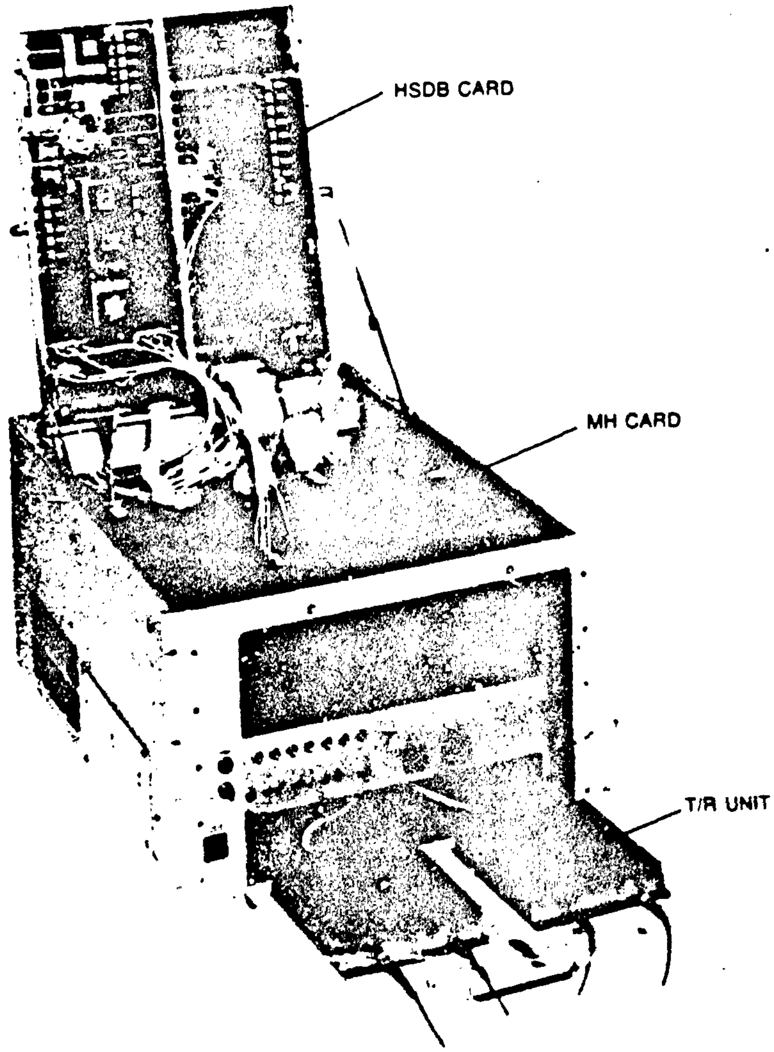

Figure 83. Brassboard Terminal

Two HSDB Tx ports provide a path through which packets generated in the terminal may be placed on the HSDB network. Tx ports are also accessible through the front panel access slot of the terminal.

The USER port provides a control/data path between the terminal and an associated Sperry 1631 processor (USER). This interface uses the PIO channel for transfer of control and status information and the DMA channel for transfer of messages to/from the HSDB network.

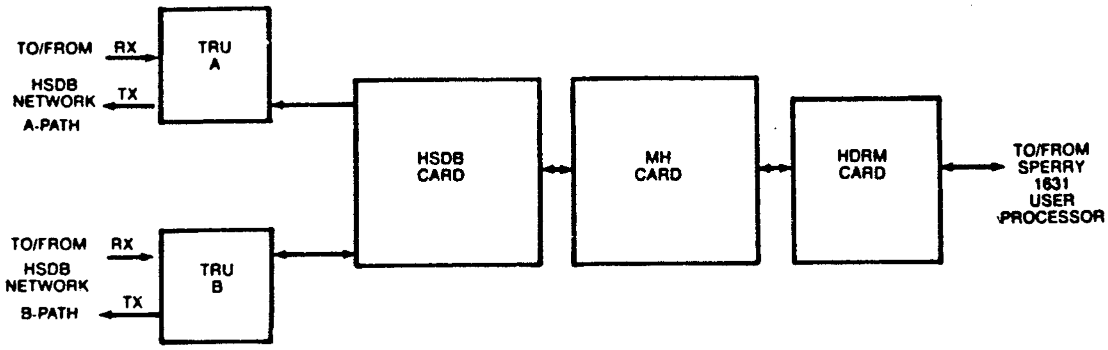

Figure 84. Brassboard Terminal Major Components

Overview of Capabilities

The brassboard terminal meets the performance characteristics for a TYPE-I network as defined in the PAVE PILLAR HSDB system specification. A bus interface unit (BIU) as described in the system specification comprises one brassboard terminal and a Sperry 1631 processor with associated peripherals (GFE). The local USER associated with the terminal is emulated by software hosted on the Sperry 1631 processor. HSDB functions are implemented using a combination of terminal hardware, terminal embedded software, terminal embedded microcode, and Sperry 1631 software. Hardware modifications to the Sperry 1631 processors or associated peripherals are not required.

- The terminal implements the entire HSDB token passing protocol defined by the HSDB system specification. The topology map is implemented for a maximum of 64 HSDB nodes. Each terminal is capable of operating as the RINGMASTER terminal.
- A reference clock function is included in the terminal. The clock is capable of operation in either TIMEKEEPER or TIMESLAVE mode.
- The terminal provides a set of timer/counter functions in accordance with the requirements of the HSDB system specification.
- The terminal contains means to accumulate the set of terminal and network statistics defined in the HSDB system specification.
- The terminal implements the CLEAR TOKEN vie process as defined in the HSDB system specification.
- BIT functions are included in the terminal which perform the following health and welfare checks of terminal operation.
  a. **Network Loopback** - The terminal monitors each of its own transmissions to verify that it is being correctly sent.

b. **Local Loopback** - The terminal is set to loopback from transmitter to receiver without the signal appearing on the network.

c. **Maintenance Loopback** - The terminal is set to loopback messages through another terminal via the network.

d. **Startup Check** - The terminal performs embedded maintenance diagnostic procedures prior to accepting the command to go ONLINE.

- The terminal implements the error recognition and recovery strategies defined in the HSDB system specification.

5.5.2.2 HSDB Card Design

Figure 85 shows functional block diagram of the HSDB card. The HSDB card is comprised of the following 4 functions:

a. Receiver/Transmitter Machine (RTXM)
b. Input Controller Unit (ICU)
c. Redundancy Management Unit (RMU)
d. State Controller Unit (SCU)

Much of the logic required was implemented using gate arrays and ancillary control/store and buffer integrated circuits. The remainder was implemented using discrete logic devices. In the following paragraphs the operational characteristics of each of the functions will be summarized. The test procedures performed on each to verify those characteristics; and the results of those tests are described in the next paragraph.

RTXM Function

The RTXM function includes the RTX Gate Array and two ringing tank cards, one for each receive channel. The ringing tank card recovers a synchronous clock signal from the received Manchester waveform.

The RTX chip is an ECL gate array designed by Rockwell with Fairchild performing as the foundry. The chip consists of two completely independent receiver sections and one transmitter section with two electrically independent transmit outputs. The chip performs the following functions:

a. Detects the start delimiter and produces a start sync signal for each message.
b. Delineates each 16 bit word and develops a word load strobe
c. Checks for and flags all Manchester and framing errors.

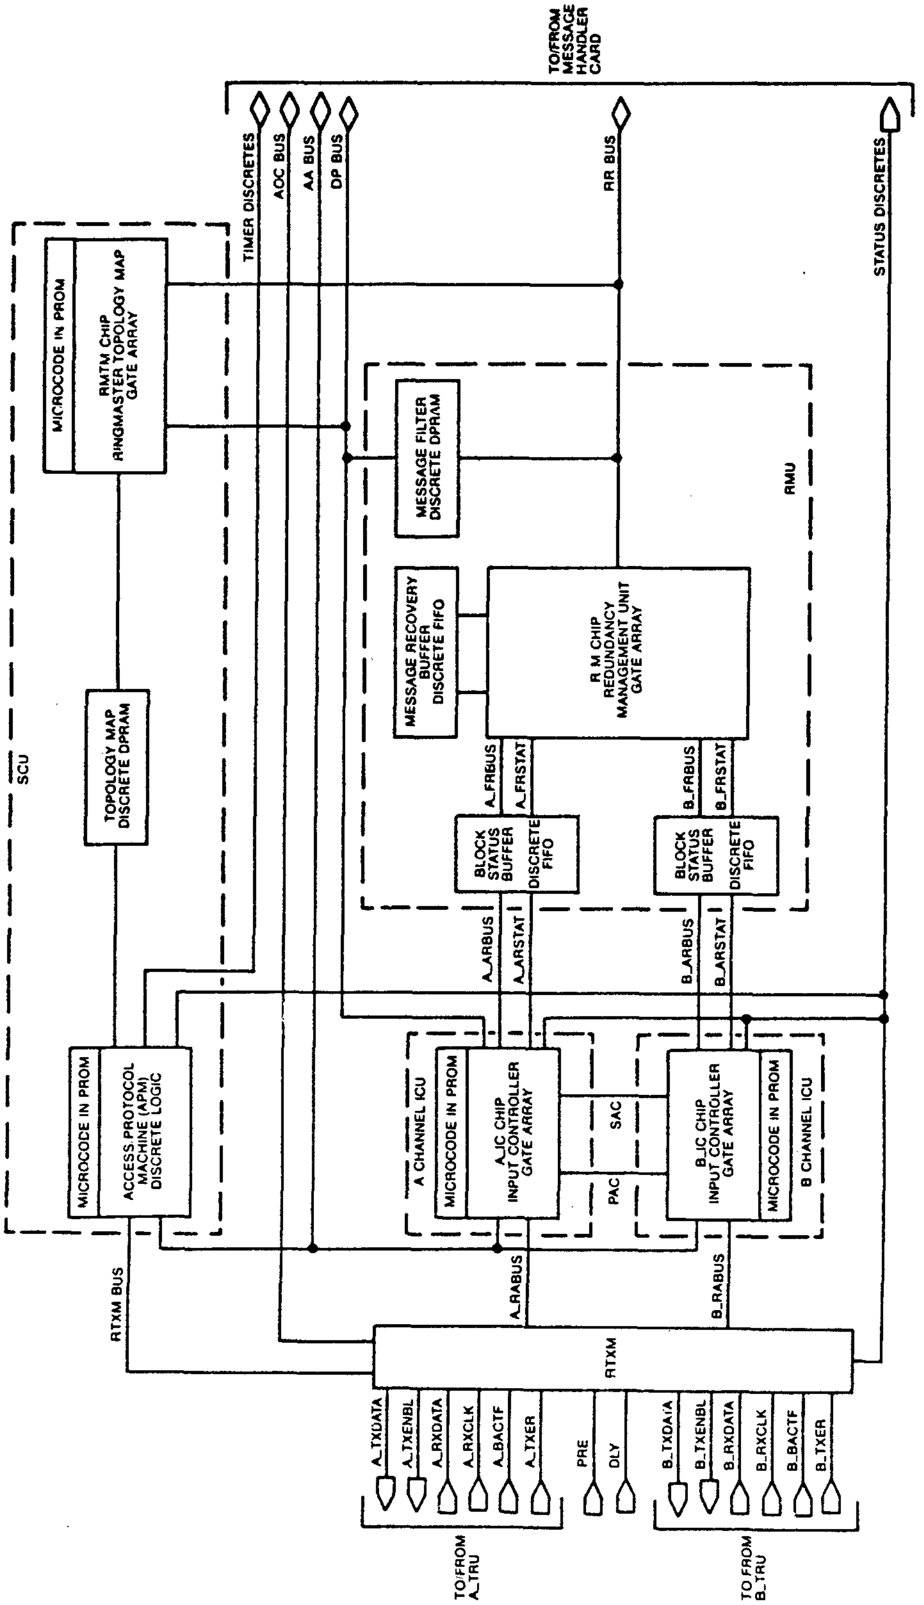

Figure 85. HSDB Card Functional Block Diagram

The transmitter section of the chip generates a Manchester encoded data signal for use by the TRU. The transmitter section of the chip also performs the following functions.

a. The chip can produce a Manchester error for BIT.
b. The chip can produce a framing error for BIT.
c. The chip can be commanded to send any length of preamble from 0 to 32 bits.
d. The chip can set the terminal response delay.
e. The chip can set to local loopback mode.
f. The chip can send the concatenate delimiter.
g. The chip can send the abort delimeter.

ICU Function

Each ICU channel is implemented with a CMOS gate array (the IC chip) which implements a custom programmable state machine. Microcode governing operation of the state machine is stored in external PROM. The ICU interfaces with the RTXM, the SCU, and the RMU. There is also an interface between the two ICU channels. Functions of the ICU include the following:

a. Receive and validate messages addressed to "this terminal".
b. Determine the destination of the data messages and direct them to either the RINGMASTER/TOPOLOGY Manager (RMTM) or to the Receive Buffer (RXB).
c. Store the current predecessor and current successor.
d. Provide the Reference Clock Time (RCT) function.

RMU Function

The RMU is the function of the terminal which coordinates the redundancy features. It consists of a three FIFO buffers, the redundancy manager (RM) CMOS gate array chip, and a dual port message filter RAM.

In operation, messages from each of the redundant ICU channels are input to a dedicated block buffer FOFI. Each block of 256 or fewer words is tagged with a status word to indicate whether it has been received without error. The RM chip retrieves the messages block-by-block and checks the associated status work; the earliest arriving correct block is written to the message recovery buffer FIFO, the redundant block is discarded. The destination address field of each message is checked against entries in the message filler RAM while being held in the message recovery buffer. Messages addressed to the terminal are forwarded to the RX buffer of the MHU, others are discarded.

SCU Function

The SCU is the functional section of the terminal that executes the protocol. It contains the logic which vies for the bus, initializes the logical ring, and operates in the token passing network. The SCU contains seven major functions:

1. The synchronization of signals between terminal functions.
2. Topology memory interface
3. Protocol State Machine
4. Discrete signal interfaces
5. BIT Error Generator
6. Arithmetic Logic Unit (ALU)
7. AA(0:15) Bus Controller

The RMTMU is a subfunction of the SCU which implements the MANAGE_TOPOLOGY_MAP process whenever the terminal is operating as Ringmaster. It accomplishes this by monitoring all network traffic, whether addressed to itself or not and maintaining a current snapshot of the network configuration.

5.5.2.3 Message Handler Card Design

The message handler card contains a data processor whose function is the management of information flow between the external user processor (Sperry 1631) and the terminal. Figure 88 shows relationship of included functions.

DPM Function

As Figure 86 indicates, the DPM supports a dual bus architecture including the following hardware functions:

a. AAMP, a Rockwell developed 16 bit microprocessor
b. 32K words of program ROM
c. 32K words of immediate data RAM,
d. An interrupt controller
e. Several counters and timers for statistics and protocol use
f. A-port to control and status registers in the protocol machine
g. User interface port

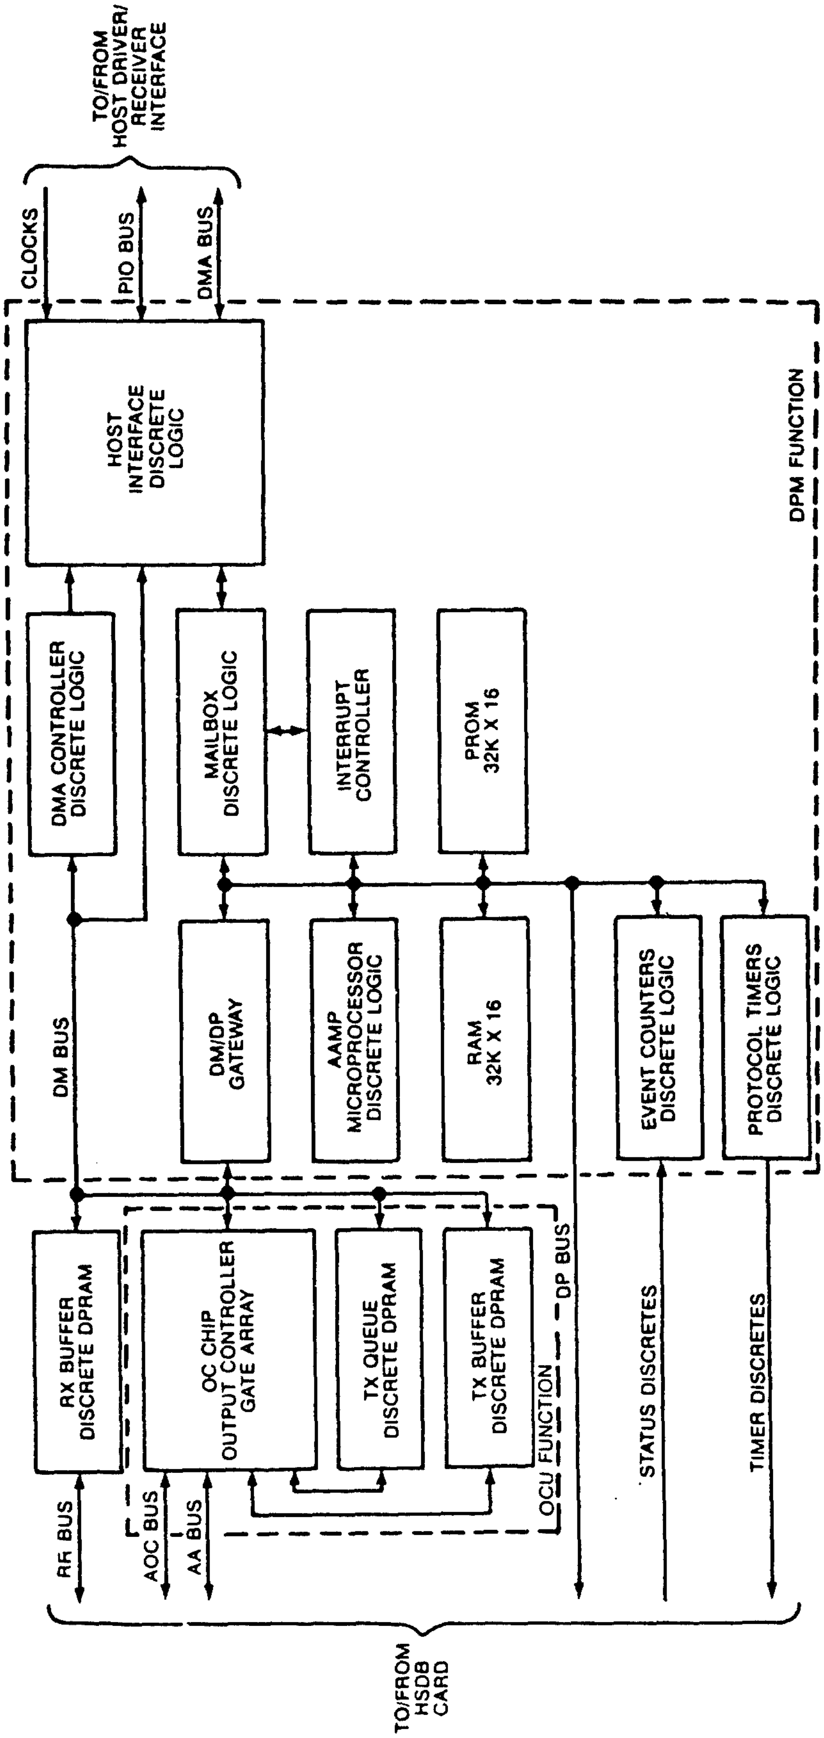

Figure 86. Message Handler Functional Block Diagram

Operating system software, included in the ROM, performs a variety of service functions for messages received from, or sent to, the user.

When a message is received from the network the SCU informs the DPM. The DPM reads the message header to determine the length, source, and priority of the message. This information is passed to the user along with the message location in the RXB. The user may recover the message at any time. A similar process is invoked for messages from the user which are to be sent on the HSDB.

OCU Function

The purpose of the OCU function is to coordinate the transmission of messages originated by the user or by the DPM function on the HSDB. As shown in Figure 86, the function consists of output controller gate array chip, the transmit queue dual port RAM and the transmit buffer dual port RAM.

When a message from the user or from the DPM is ready for transmission, it is placed in the Tx Buffer. When entry has been completed, a transmission request is placed in the Tx queue. This prompts the output controller to schedule and send the message in accordance with the protocol.

5.5.2.4 TRU

Each TRU is a metal-module assembly which contains the five circuit boards required for a single fiber optic transmitter channel and a single fiber optic receiver channel. This method of packaging was selected because of the need for extensive shielding for the fiber optic receiver circuits. Figure 87 shows a block diagram of the TRU. The design is straight forward and based on the design of the Task II TRU previously described so that detailed design information is not included.

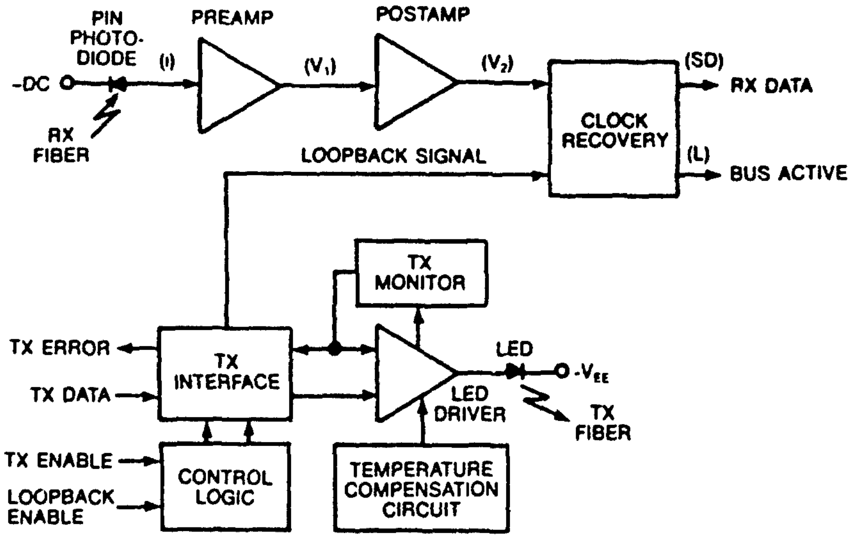

Figure 87. The TRU Was Updated From The Task II Design To Make It More Producible

6.0 TESTING, CHARACTERIZATION AND DEMONSTRATION

Designs for the HSDB were validated using a combination of testing, characterization and demonstration.

Testing

All critical circuits were breadboarded and tested prior to being integrated into the breadboard and brassboard designs. Testing was done on an informal basis with no documentation provided as CDRL deliverables.

Characterization

TRU hardware operation was characterized as part of Task I/II effort. Acceptance Test Plans governed each characterization and test reports were prepared at the culmination of each. The result of coaxial TRU and Fiber Optic TRU characterization were presented at Task I and Task II ATR, respectively.

Figure 88 shows the characterization test equipment developed for the program as it was configured for coaxial TRU characterization. It uses a HP9826 programmable calculator to control a HP8180 data generator, a HP8182 data analyzer, a Fluke 1953A frequency counter timer, and a Micronetics PNG 5017 noise generator. The special test equipment panel located in the center of the test rack and the card cages located on the work surface were used to interface with the TRU being tested.

Demonstration

Four demonstrations were conducted as part of the program:

1. A coaxial HSDB demonstration was held as part of the Task I ATR.
2. A fiber optic HSDB demonstration was held in conjunction with Task II ATR.
3. A core protocol demonstration was held in conjunction with the Task IV interim design review.
4. A PAVE PILLAR protocol demonstration was held in conjunction with the Task IV final demonstration review.

Demonstrations (1) and (2) utilized the characterization test equipment described above. Demonstrations (3) and (4) used the HSDB demonstration system fabricated specifically for that purpose. Figure 62 shows the HSDB demonstration system as it was configured for the breadboard protocol demonstration. This test bed functioned as the principal vehicle for

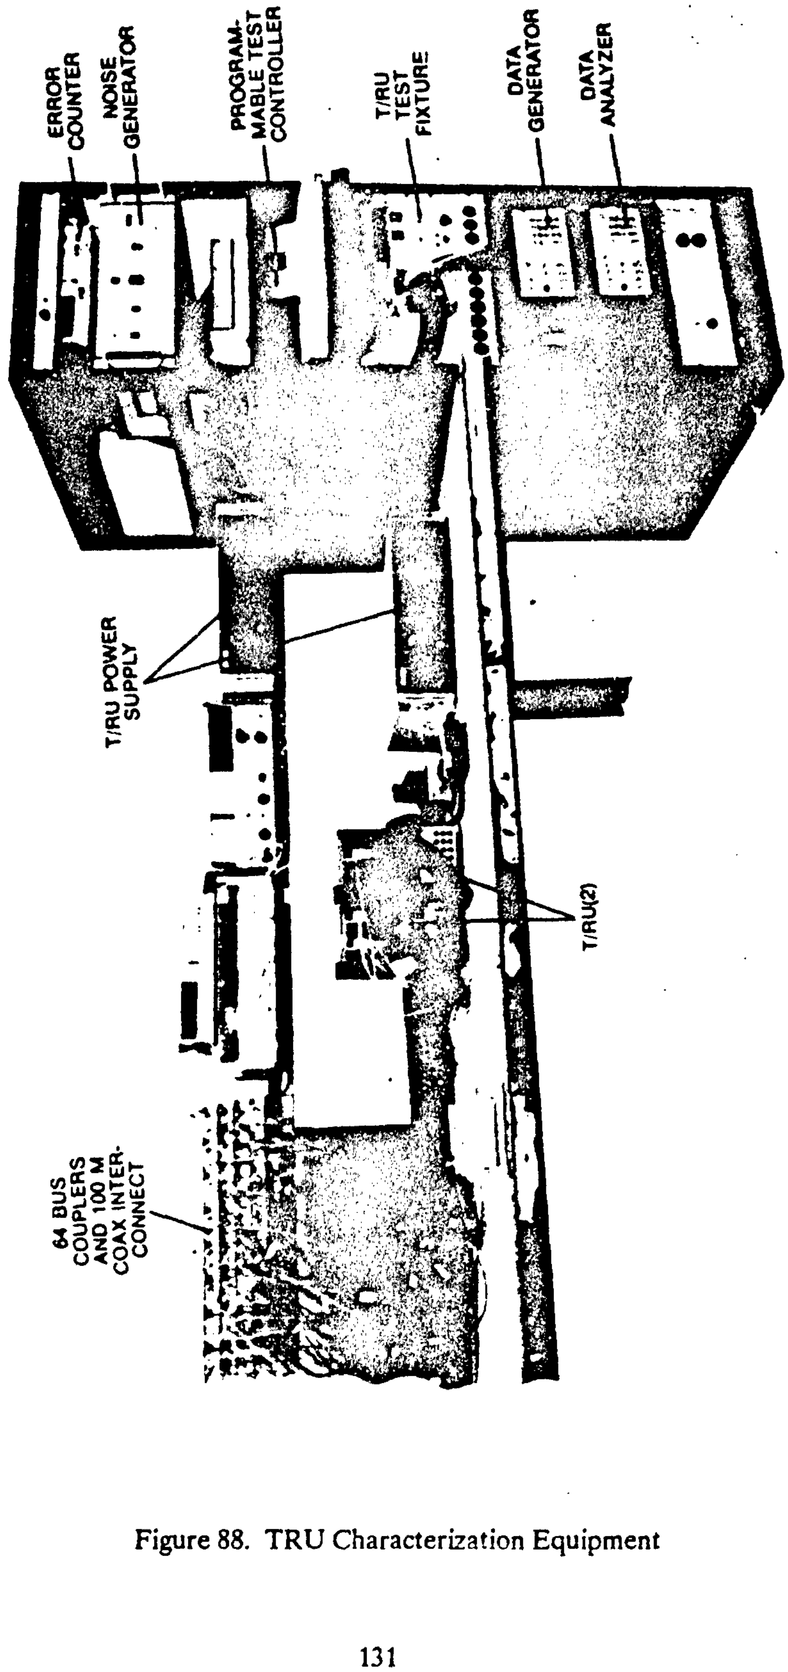

Figure 88. TRU Characterization Equipment

verification/validation of the protocol design as well as providing demonstration hardware. The design of the HSDB terminal used for verification/validation demonstration, including the semi-custom chip set developed for this program, is described in paragraph 6.4.

This series of tests, characterizations, and demonstrations has verified the practicality of a 50 Mbps HSDB. While hardware developed for these validation activities is not directly applicable to production systems, it does show that production hardware with acceptable performance, can be developed at an acceptable cost and risk.

6.1 Validation of Coaxial HSDB Characteristics

Performance characteristics of the brassboard coaxial TRU and the coupler was determined by testing the following:

a. Transmitter output power
b. Transmitter output waveform
c. Transmitter clock stability
d. Transmitter synchronization waveform
e. Transmitter timeout override
f. Transmitter switch waveform
g. Transmitter output noise
h. Manchester encoding
i. Receiver acquisition range
j. Receiver lock time
k. Receiver dynamic range
m. Receiver input impedance
n. Bit-error rate
o. Coupler mainbus
p. Coupler Tx port
q. Coupler Rx port

Bit error rate performance of the TRU was characterized over the temperature range of -54 ^{o}C to +95 ^{o}C.

Transmitter Output Power

Each transmitter was tested to verify that it would produce an output between +1.25 Vp-p and 1.7 Vp-p when loaded by 50 Ohms. The six brassboards produced outputs between 1.55 Vp-p and 1.67 Vp-p.

Transmitter Output Waveform

The output waveform of each transmitter was tested against the requirements shown in Figure 89. All were found to be well within the acceptable range of operation.

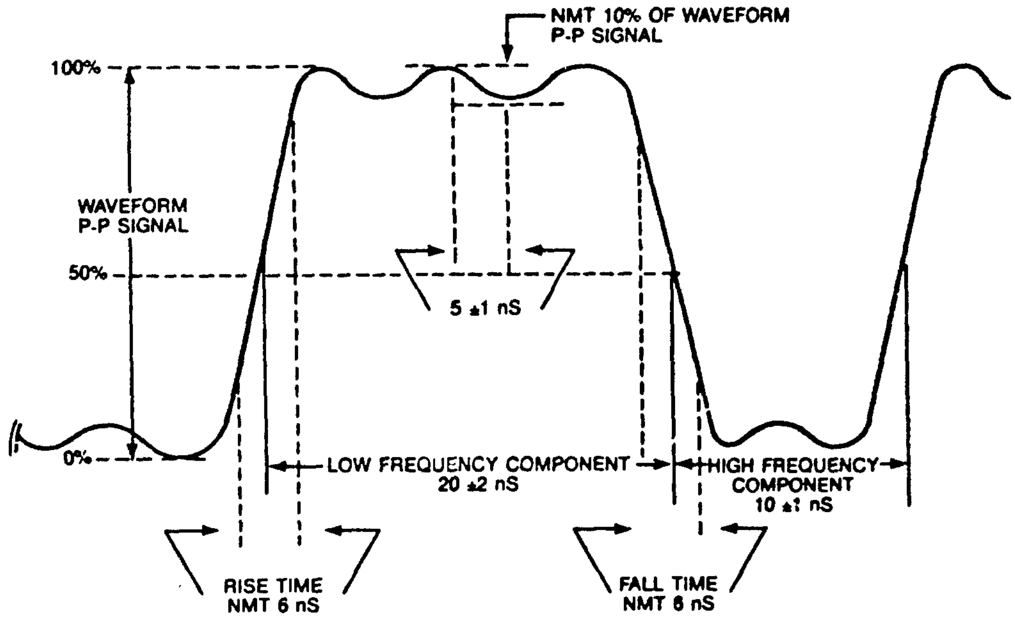

Figure 89. Coaxial Transmitter Output Waveform Requirement

Transmitter Clock Stability

The clock of each transmitter was measured to verify that its center frequency was 50 MHz + 12.5 kHz and that frequency components removed greater than 12.5 kHz from the center frequency were at least 40 dB below the level of the carrier.

Synchronization Waveforms

Preamble, start delimiter and end delimiter waveforms from each transmitter were verified against the following criteria:

a. Preamble: 8 bits of logical 0 bits
b. Start delimiter: per Figure 95(a)
c. End delimiter: per Figure 95(b)

All were found to produce correct synchronization waveforms. (See Figure 90)

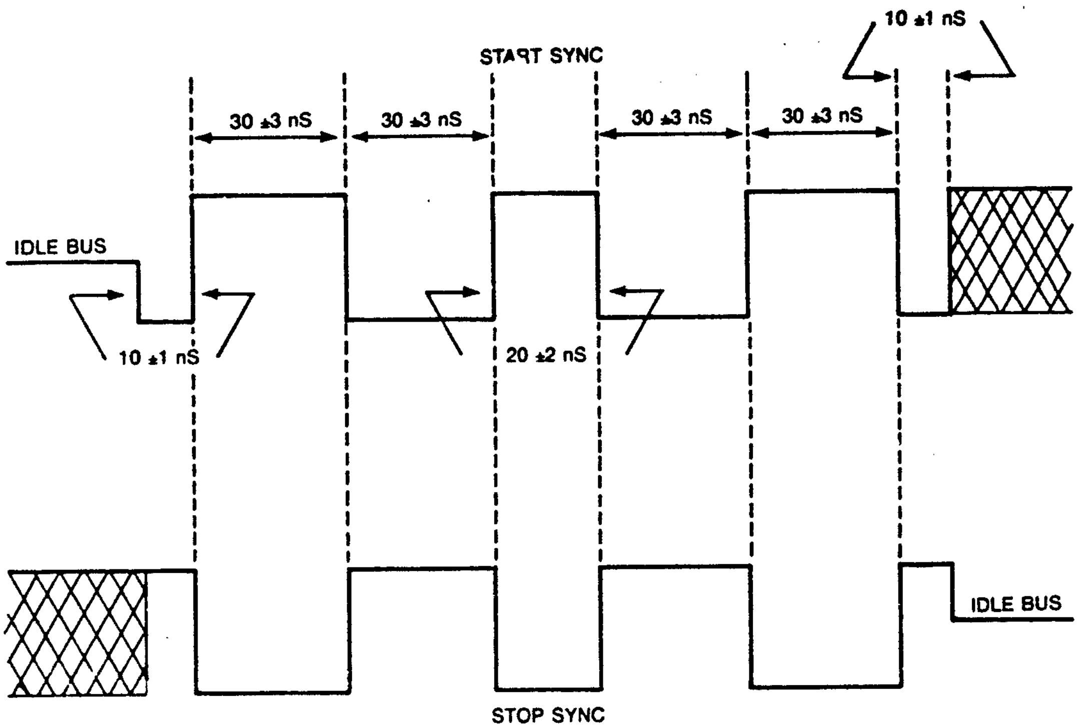

Figure 90. Synchronization Waveforms are Illegal Manchester Symbols

Transmitter Timeout Override

The watchdog timer of each transmitter was tested to verify that it would limit transmissions to NMT 1.44 mS. The six brassboard transmitters terminated transmission at between 1.20 mS and 1.34 mS.

Transmitter Switch Waveform

Each transmitter was tested to verify that it would produce a square switch control signal at the switch control port which enclosed the packet being sent from the output port. All TRUs produced an acceptable waveform.

Transmitter Output Noise

The noise produced at the output port of each transmitter during receive operation was measured using a broadband power meter. The highest level measured was -18dBm against a limit of -12 dBm.

Manchester Encoding

Each transmitter was tested to prove generation of valid Manchester waveforms when driven using a pseudo random serial bit stream. The waveforms were shown to have transitions at 10 ±1 nS and 20 ±2 nS at half-cycle points.

Rcceiver Acquisition Range

The acquisition range of each receiver was tested using a transmitter modified to allow the carrier frequency to be varied above and below 50 Mhz. All receivers met their bit error rate specification while the transmitter carrier frequency was modulated at 1000 Hz with a deviation of +25 KHz.

Receiver Lock Time

Each receiver was tested to prove that it acquired synchronization within 4 preamble bit-times.

Receiver Dynamic Range

Each receiver was tested for dynamic range performance using the test setup shown in Figure 91. Packets of pseudo random data were sent alternately from transmitter #1 (Tx1) and transmitter #2 (Tx2) to the receiver under test. Network components were adjusted so that alternate packets were different in amplitude, Tx1 being 250 mVp-p and Tx2 being 22 mVp-p. All receivers met their bit error rate specification while receiving these alternate high level/low level packets.

Receiver Input Impedance

The input impedance of each receiver was shown to be within the range 50 ±10 Ohms at 0 ±10 degrees between 10 MHz and 60 MHz.

Bit-Error Rate

Bit-error rate testing was performed on each TRU. No attempt was made to allocate errors against either the transmitter or the receiver although it appeared as if essentially all errors were receiver-induced. Figure 92 shows the test setup used. Tx1 and Tx2 were set to produce alternate packets at receiver operating range limits as described above for the dynamic range test. Broadband noise was injected into the network to produce the desired S/N ratio on the low level packet at the input port of the receiver. All receivers met their specification requirement at room temperature. Two were tested across the temperature range -54 ^{0}C to +95 ^{0}C. The results of this test are shown in Figure 93. One of the two met its requirement across the full temperature range. The second was outside the range at cold temperatures. Bench

testing performed after the characterization showed that the clock recovery circuit was slightly misadjusted, causing the problem.

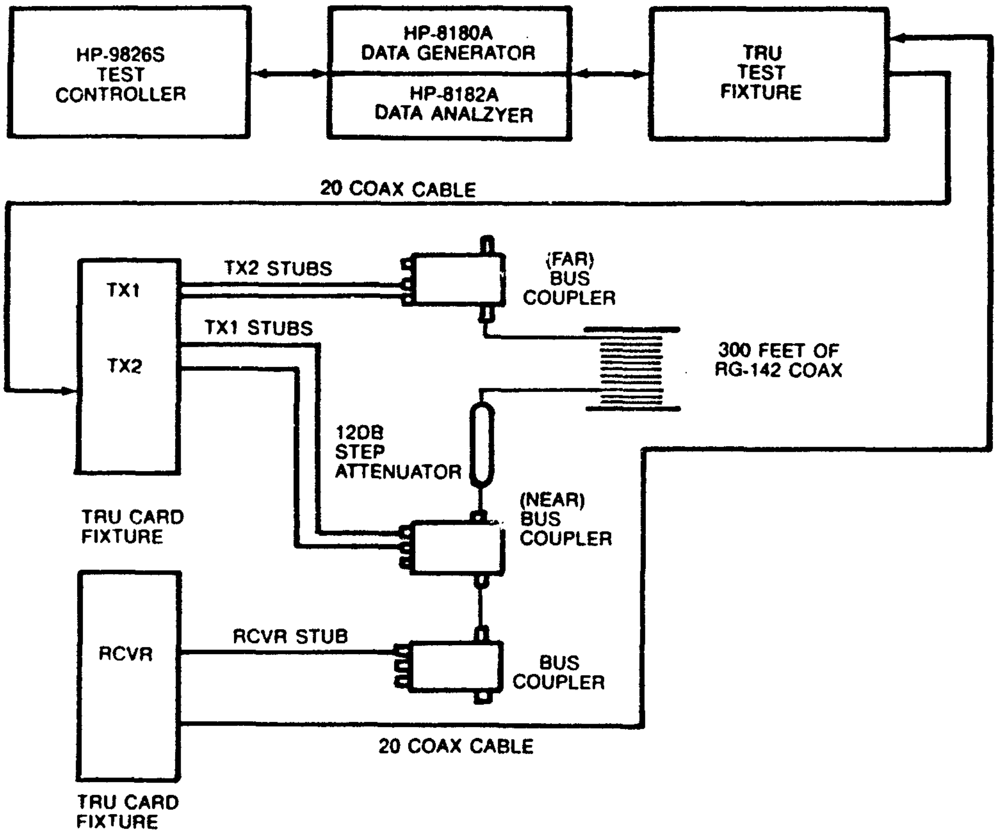

Figure 91. Receiver Dynamic Range Test Setup

Coupler Mainbus

The mainbus ports of each prototype coupler were tested for loss, return loss, and group delay. The highest loss measured was 0.12 dB, well within the requirement of 0.2 dB. Return loss met the requirement of NLT 32 dB in all cases. Group delay was measured at 0.6 ns worst case. These measurements proved the practicality and state of the art performance of the coaxial HSDB.

Coupler Tx Port

The Tx port of each coupler was tested for coupling coefficient, isolation and group delay. Coupling measured between 10.0 dB and 10.2 dB at 25 MHz and between 10.9 dB and 11.3 dB at 50 MHz. Isolation measured greater than 52 dB; group delay was 5 ns for all four prototypes.

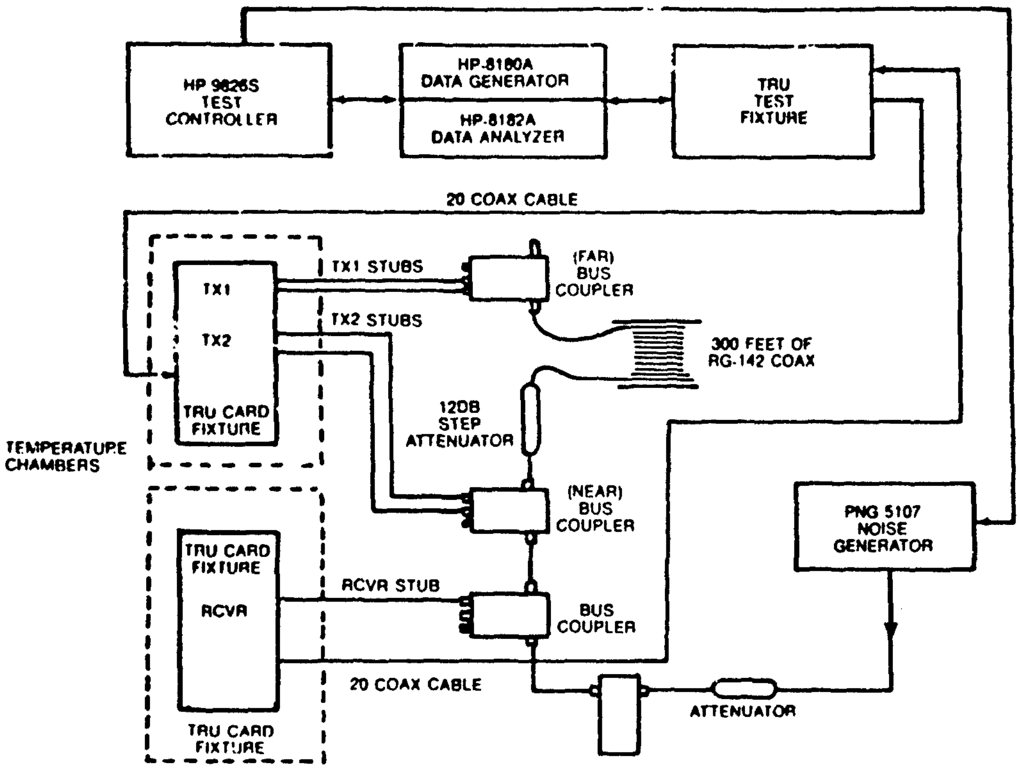

Figure 92. TRU Bit-Error Rate Test Setup

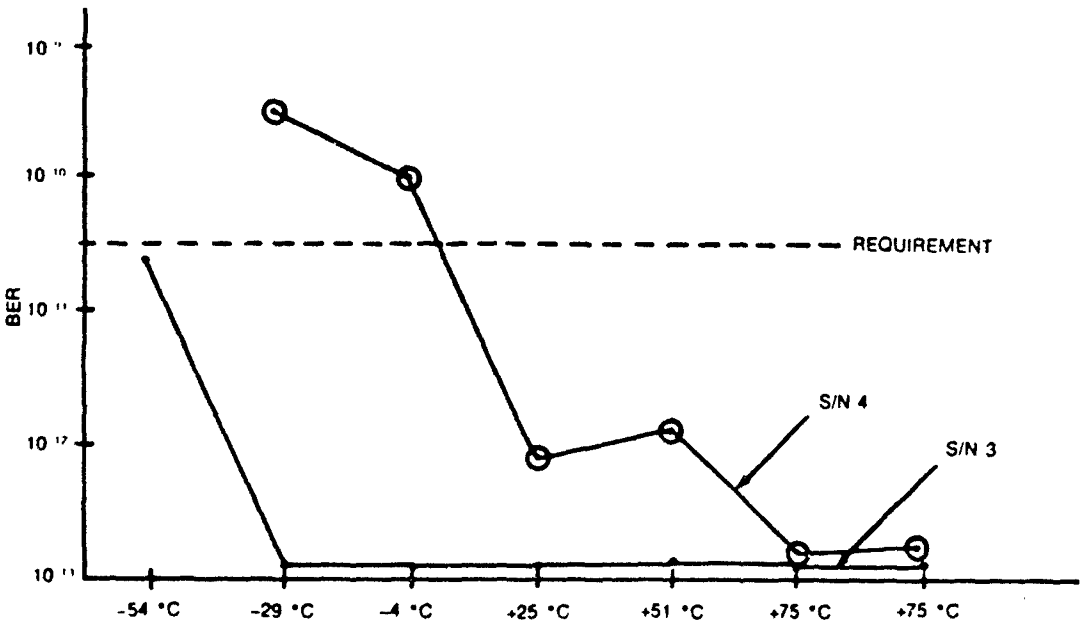

Figure 93. Bit-Error Rate Characterization Results

Coupler Rx Port

The Rx port of each coupler was tested for coupling coefficient and group delay. Coupling measured between 20.0 dB and 20.4 dB at 25 MHz, and between 19.8 dB and 20.4 dB at 50 MHz. Worst case deviation between 25 MHz and 50 MHz for any single coupler was 0.2 dB. Group delay measured between 2.0 nS and 2.4 nS for the prototypes.

6.2 Validation of Fiber Optic HSDB Characteristics

Performance characteristics of the brassboard fiber optic TRU was determined by testing the following:

a. Transmitter Power Output
b. Transmitter Waveform
c. Transmitter Clock Stability
d. Transmitter Synchronization Waveforms
e. Transmitter Timeout Override
f. Manchester Encoding
g. Receiver Acquisition Range
h. Receiver Dynamic Range
i. Bit-Error Rate
j. Interpacket Gap

Bit error rate performance of the TRU was characterized over the temperature range from -54 °C to +95 °C.

Transmitter Output Power

Each transmitter was tested to verify that it would produce a peak output power between -4 dBm and -7 dBm. The six brassboard produced outputs between -4.9 dBm and -6.2 dBm.

Transmitter Waveform

The output waveform of each transmitter was tested against the requirement shown in Figure 94. All were found to be within the acceptable range of operation.

Transmitter Clock Stability

The clock of each transmitter was measured to verify that its center frequency was 50 MHz + 12.5 KHz and that frequency components removed greater than 12.5 KHz from the center frequency were at least 40 dB below the level of the carrier.

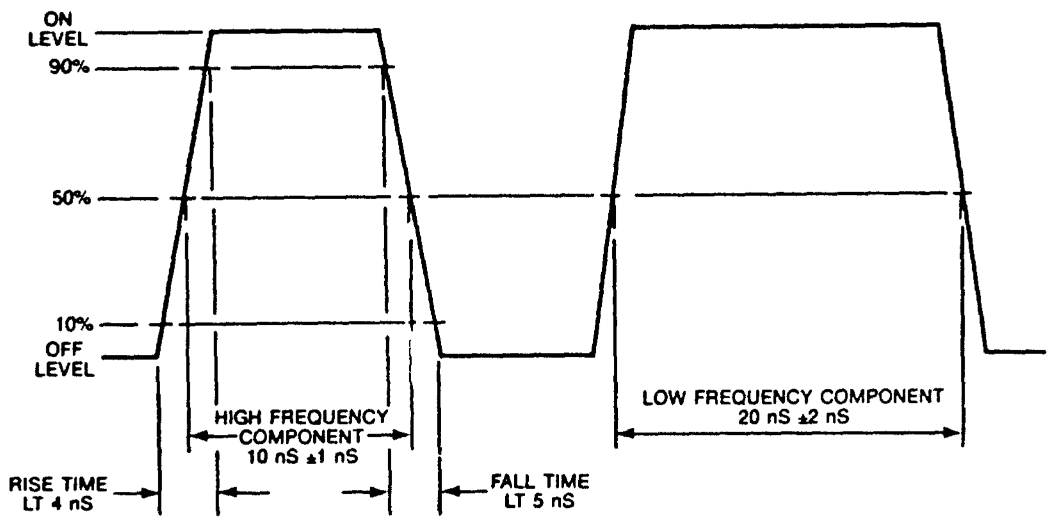

Figure 94. Fiber Optic Transmitter Waveform Requirement

Transmitter Synchronization Waveforms

Preamble, start delimiter and end delimiter waveforms from each transmitter were verified against the following criteria:

a. Preamble: 14 bits of logical 0 bits

b. Start Delimiter: Per Figure 95(a)

c. End Delimiter: Per Figure 95(b)

All transmitters were found to produce correct synchronization waveforms.

Transmitter Timeout Override

The watchdog timer of each transmitter was tested to verify that it would limit transmissions to NMT 1.6 mS. The four brassboard transmitters terminate transmission at between 2.0 and 2.2 mS. This is an area which obviously requires redesign before production but the deficiency did not cause problems with the remaining characterization and demonstration so no redesign was initiated as part of Task II.

Manchester Encoding

Each transmitter was tested to prove generation of valid Manchester waveforms when driven a pseudo random serial bit stream. The waveforms were shown to have transitions at 10 ±1 nS and 20 ±2 mS at half peak power points.

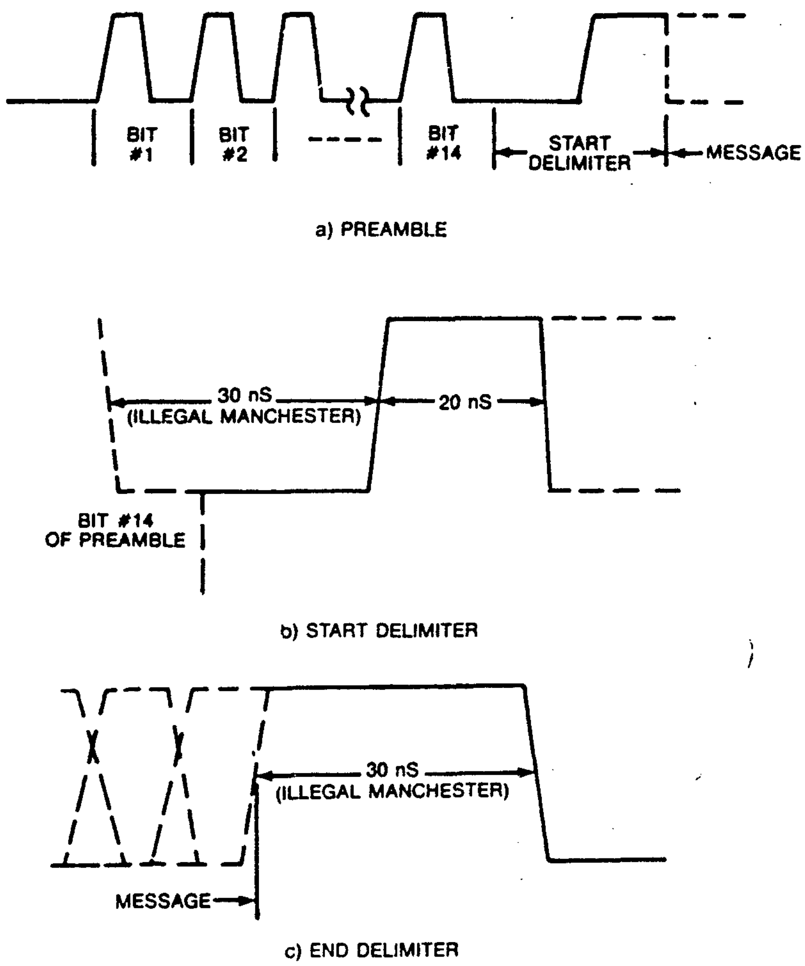

Figure 95. Fiber Optic Synchronization Waveforms for Task II TRU

Receiver Acquisition Range

The acquisition range of each receiver was tested using a transmitter modified to allow the carrier frequency to be varied above and below 50 MHz. All receivers met their bit-error rate specification while the transmitter carrier frequency was modulated at 1000 Hz with a deviation of +25 KHz.

Receiver Dynamic Range

Each receiver was tested for dynamic range performance using the test setup shown in Figure 96. Note that this does not provide operation over the full receiver operating range since there is no way to adjust the high level transmitter to a level which causes errors at the receiver. It does, however, give a good indication of receiver performance in a fiber optic HSDB network.

Differential path losses in a star topology network are limited by the nature of the interconnect and are much less severe than typically found in a linear bus network.

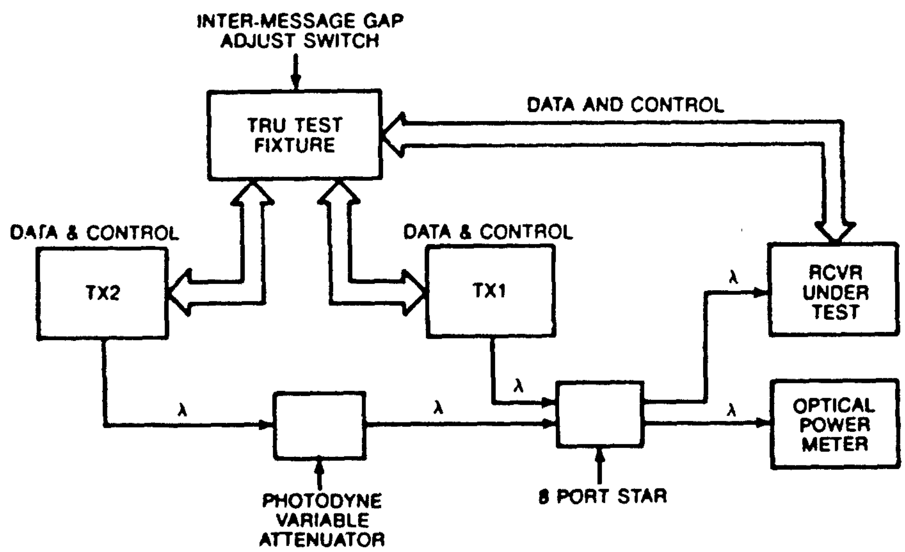

Figure 96. Fiber Optic Receiver Dynamic Range Characterization Setup

Packets of pseudo random data were sent alternately from transmitter #1 (Tx1) and transmitter #2 (Tx2) to the receiver under test. Network components were adjusted so that alternate packet were different in amplitude, Tx1 being -15 dBm peak and Tx2 being -27 dBm peak. All receivers operated adequately for this test.

At the Task II ATR questions were raised concerning the temperature stability of the receivers. As a result, Rockwell performed additional characterization testing in this area. Six brassboard receivers were subjected to temperature extremes from -54 ^{o}C to +90 ^{o}C. The receiver input signal for each was varied to the level at which no errors were observed over a 2 minute period of operation. Figure 97 shows the results of that testing. This shows that the design functions over the target temperature range with fairly stable characteristics.

Bit-Error Rate

BER testing was performed on each TRU. No attempt was made to allocate errors against either the transmitter or the receiver although it appears as if essentially all errors were receiver-induced. The bit-error rate test was conducted under identical conditions described above. Note that this test was not carried out at the receiver sensitivity design point of -37 dBm peak. Preliminary testing showed this goal to be unattainable. Instead -32 dBm appeared to be a typical room temperature sensitivity. A derating of 2 dB for temperature range variation and 3

dB for component variation were added in accordance with standard engineering practices, to arrive at the test specified sensitivity operating point of -27 dBm. All TRUs met this specification.

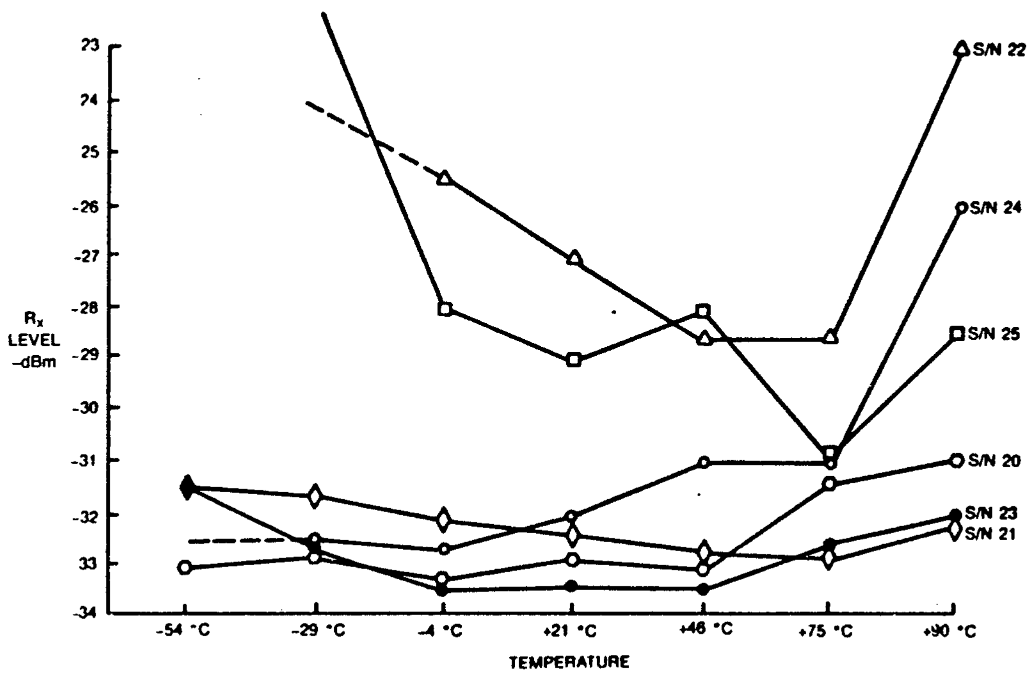

Figure 97. Fiber Optic Receiver Sensitivity vs. Temperature

Interpacket Gap

At the Task II ATR questions were raised as to the impact of gap time between packets on receiver dynamic range performance. As a result, Rockwell performed additional characterization testing in this area. The test conditions used for dynamic range characterization were repeated with the exception that the interpacket gap was varied from 50 ns outwards while the dynamic range was similarly varied from 0 to 20 dB. Full test data is provided in the test report. To summarize most receivers would operate with 50 ns interpacket gap time for dynamic range conditions of less than 10 dB. Required gap time increased to several μs under high dynamic range conditions. This characteristic was traced to to Q of the ringing tank oscillator used for clock recovery in the brassboard receivers. Careful adjustment of the receivers showed that all could be made to operate under conditions of wide dynamic range plus narrow gap time. Production receiver designs will need a lower Q tank circuit if acceptable reliability performance is to be achieved.

6.3 Validation of HSDB Protocol Characteristics

HSDB performance characteristics ascribable to the protocol were validated by construction, test, and demonstration of three terminals in the HSDB demonstration system. Two generations of prototype hardware were fabricated. The breadboard demonstration implemented core protocol functions of Version 3.0 of the PAVE PILLAR HSDB system specification. This demonstrated:

a. Establishment of the logical ring from a cold start
b. Transmission and reception of messages
c. Token passing
d. Recovery from a lost token
e. Adding nodes to an operating network
f. Critical terminal timing

The brassboard was designed to implement the full functionality described in Version 4.2 of the PAVE PILLAR HSDB system specification. This includes those functions described above and also:

a. 4-level priority system on messages
b. Reference clock timer
c. Topology map
d. Redundancy
e. Statistics

Because of program constraints only a single brassboard terminal was fabricated. This did not allow demonstration of an operating network but a majority of the Version 4.2 functions were validated through the use of the E1 (engineering model-serial number 1) brassboard terminal connected to a data generator/analyzer and other laboratory test equipment as shown in Figure 98.

- The HP8180A, under control of the HP9826S, was used to simulate messages arriving from the network.
- The AAMP development system provided a debugging 'window' into the DPM. The user interface was simulated and terminal embedded functions were initiated and monitored.
- The HP1630A Logic Analyzer was used to monitor a variety of points internal to the brassboard. Timing was accurately measured, flags monitored, etc.

While this approach is not ideal from the perspective of accurately mimicking an actual HSDB network, it proved to be a successful method of validating protocol operation. Specific validation scenarios are described in subsequent paragraphs.

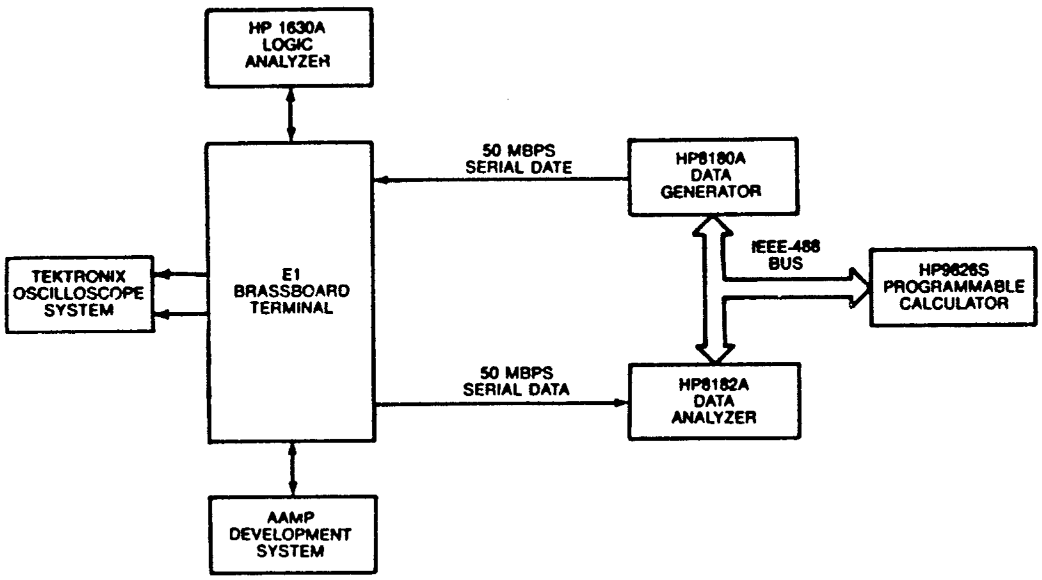

Figure 98. E1 Brassboard Terminal Test Setup

The two-phase (breadboard/brassboard) development approach was chosen as the lowest risk commensurate with the maturity of the governing system specification. The breadboard provided the vehicle for proving basic protocol functions and for validating critical timing of terminal functions. Its construction, using discrete parts mounted on plug-in circuit boards, would not allow development of a full-function prototype, however. The critical timing involved in the terminal could not be accommodated within the large format of the breadboard. Also, there was just not enough room to package all of the required electronics.

The brassboard built upon the design of the breadboard by reducing those circuits already validated into gate array designs. This allowed approximately a 10:1 reduction in physical size and allowed the full-function brassboard terminal to be packaged in the same volume as the core function breadboard. Five gate array designs are included in the PAVE PILLAR HSDB set. These are listed in Table 21. These five semi-custom devices, augmented by a variety of memory devices, provide most of the logic required to implement a PAVE PILLAR HSDB terminal. The chip set along with the required ancillary chips is shown in Figure 99. The two gate arrays along the left are the OC and the RM, the two along the right are ICs, the one in the lower center is the RTX, and the one in the upper center is the RMT.

The protocol controller unit (PLU), of the breadboard was not miniaturized at the same time as the other functions due to ongoing evolution of the protocol.

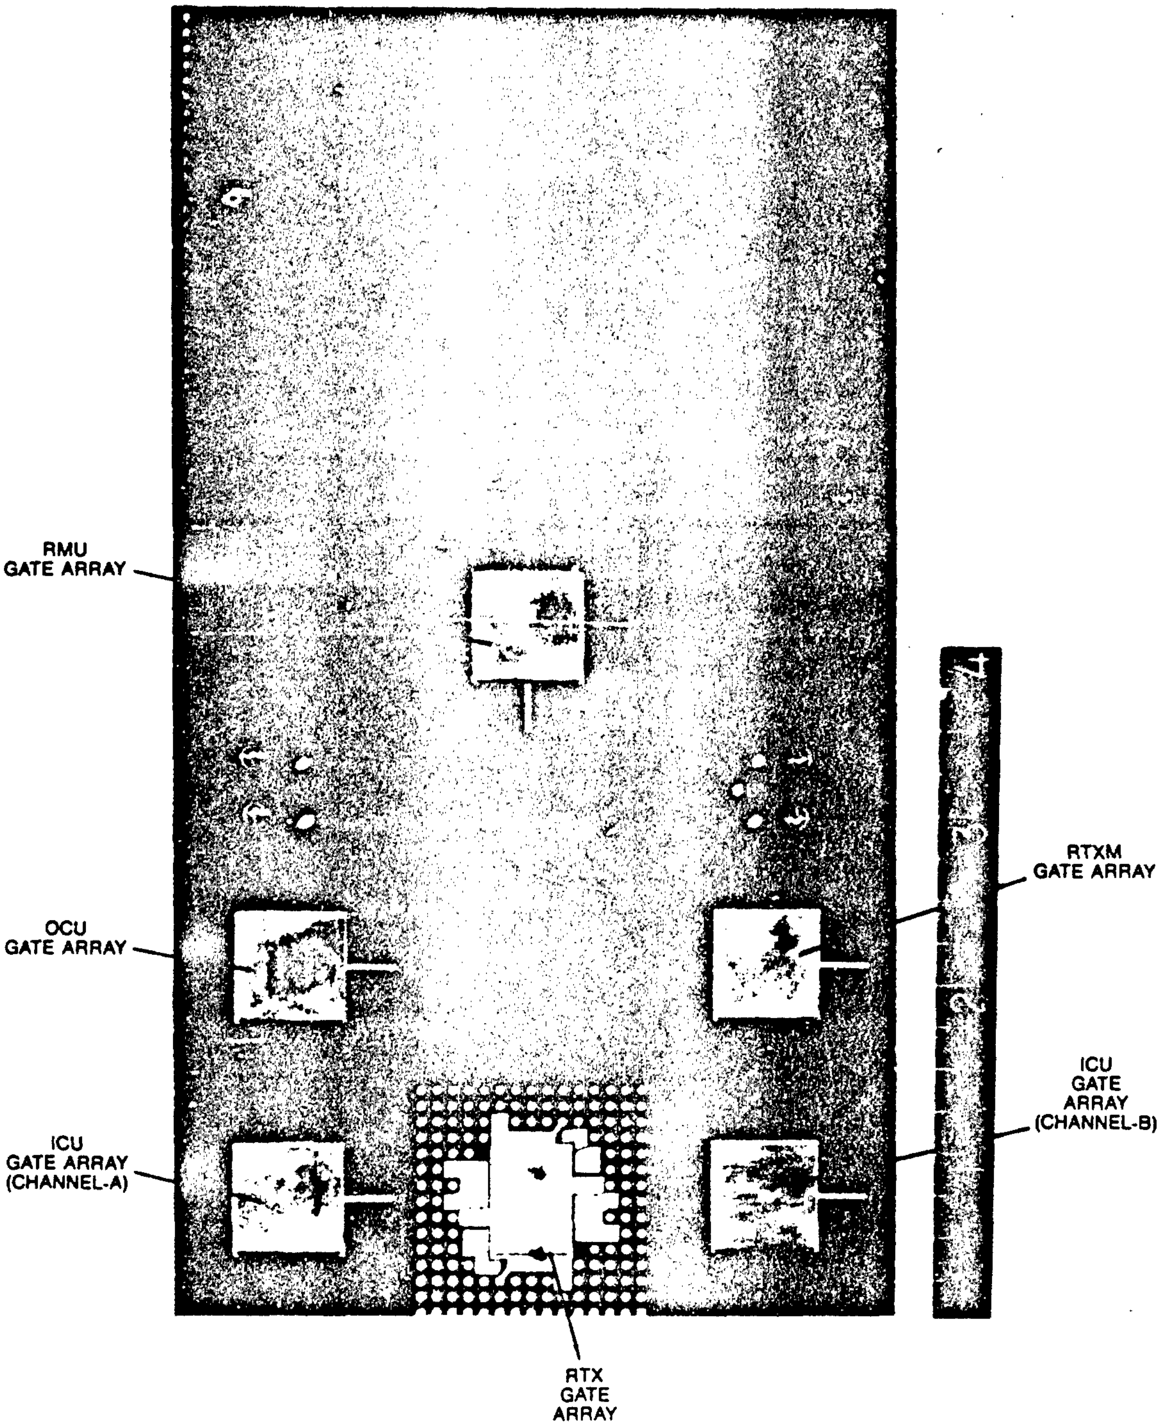

Figure 99. Brassboard Semi-Custom Chip Set with Additional Support Chips

Table 21. Brassboard Semi-Custom Chip Set

RTX -	ECL gate array contains high speed front end logic for 1 transmit channels and 2 receive channels.
IC -	CMOS gate array contains the logic for 1 receive channel (2 required).
OC -	CMOS gate array contains the logic for the dual transmit channel.
RM -	CMOS gate array contains redundancy logic for the receive channel
RMT -	CMOS gate array contains logic for the RINGMASTER and topology map manager functions.

6.3.1 Message Transmission and Reception

Verification/validation of the ability to transmit and receive messages was accomplished in two phases. The breadboard terminal included provisions for sending and receiving messages of single-level priority. The multi-level priority system, including token rotation timers, was validated with the brassboard terminal.

Figure 100 shows a packet comprised of a 20 word message concatenated with a token. This photograph was produced by connecting an optical-to-electrical converter to an unused output port of the star coupler, and applying the electrical output to the vertical channel of an oscilloscope. Notice that the low frequency component of the waveform is slightly higher in amplitude than the high frequency component. This is attributable to the rise/fall time of the LED in the transmitter.

Figure 101 also shows the concatenation delimiter between messages, in this case a token appended to a message. Note that the delimiter pair requires 8 bit times and does not contain idle time or preamble bits. This demonstrates the concatenation requirement of the system specification.

HSDB Card Testing

The transmit and receive operational characteristics of the HSDB Card were tested using a Hewlett Packard 8180 data generator in conjunction with breadboarded special test circuits mounted on cards located in the modified chassis of a Phase III BIU. Figure 102 shows the test setup. In the first test configuration the test circuits were used to supply test signals to the A and B Channels of the HSDB Card:

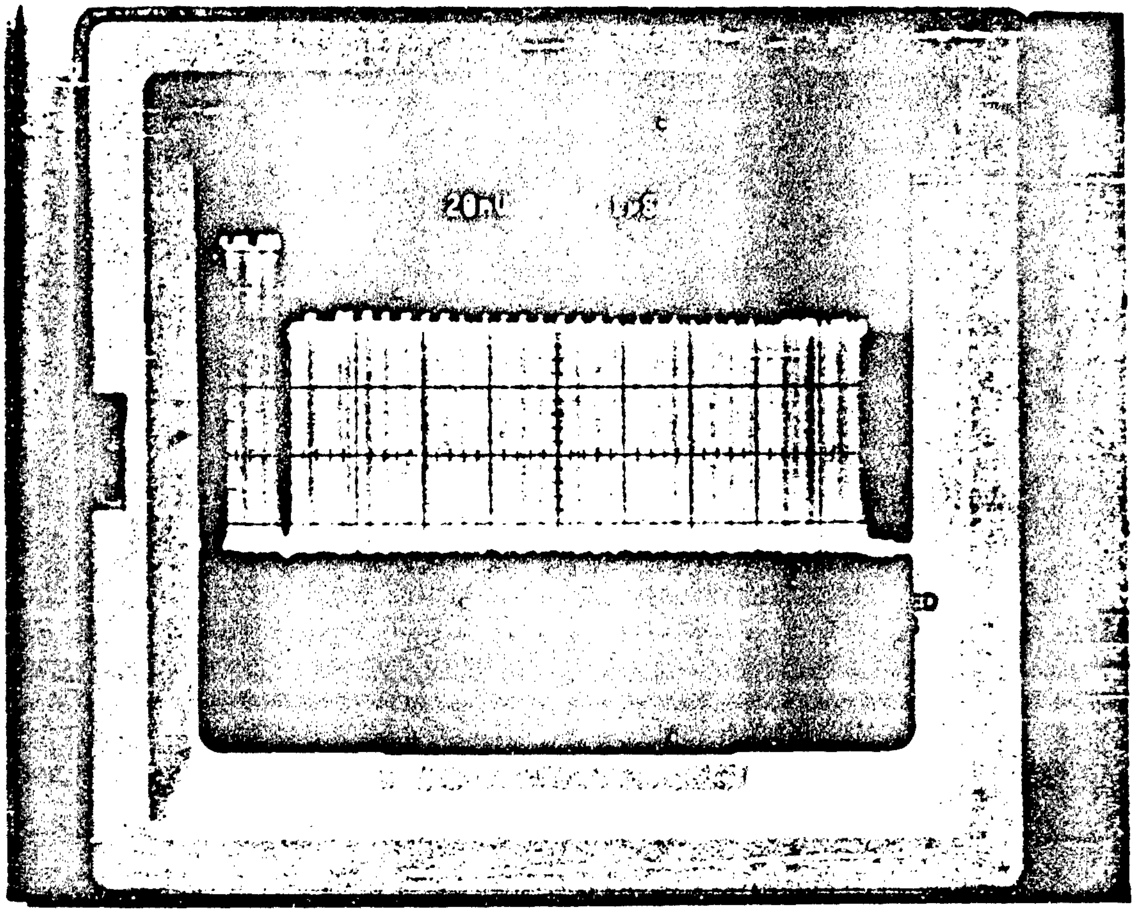

Figure 100. 20-Word Message with Concatenated Token

TEST 1: Channel A signals could be routed through any of 3 paths: (a) Directly (no relative delay), (b) Through a 150 ft. fiber optic path which created a 240 nS delay, and (c) Switching between the direct route and the 240 nS delay route on alternate transmissions. Channel B was always routed through 75 ft. of fiber optic cable. This resulted in a fixed 120 nS delay. Since Channel A could be made to alternate between relative transmission delays of 0 ns and 240 nS and Channel B has a fixed transmission delay of 120 nS Channel A (and B) will alternate between being received first and last. This test was used to verify the redundant bus selection logic of the brassboard receiver.

TEST 2: Data transmissions could be cut off midstream (no ED or ABORT) on the data message being transmitted on one channel. This test feature created a situation in which the terminal begins to receive Channel A first and then must respond by switching to and recovering the delayed Channel B message.

TEST 3: Manchester data transmissions clock frequencies were varied to test the performance of the synchronization ability of the receiver.

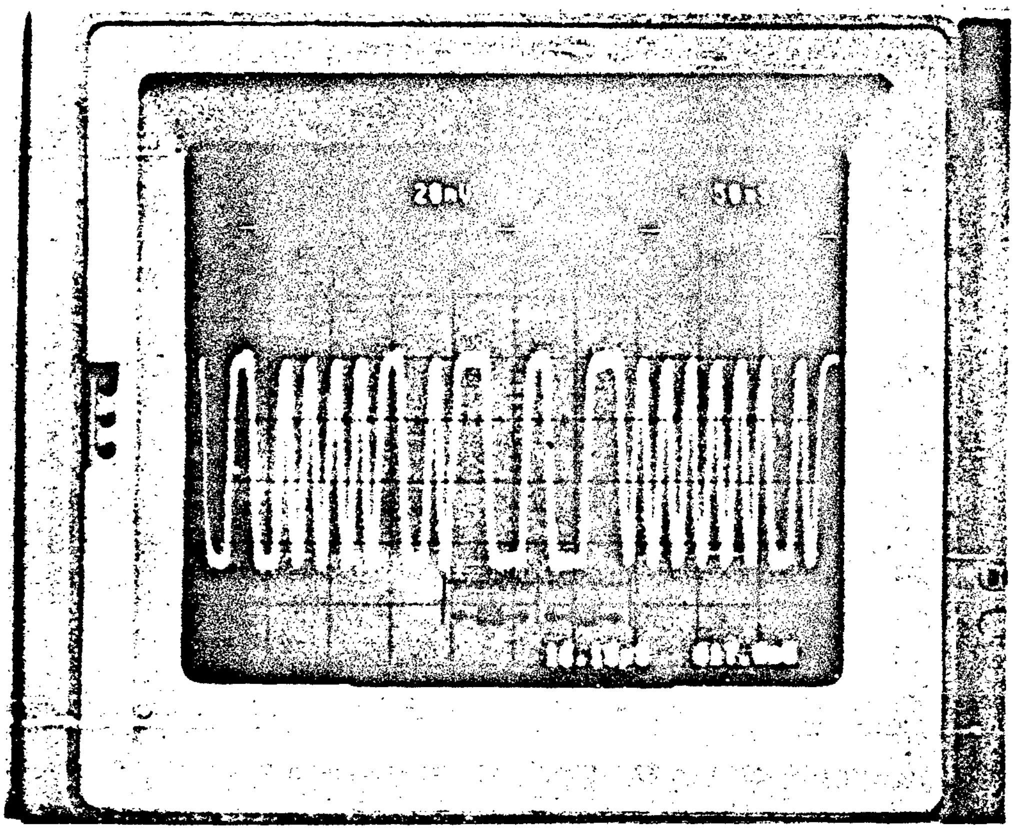

Figure 101. The Concatenation Delimiter ED/SD Separates Messages

TEST 4: The HSDB Card's RTXM transmit section was commanded to generate specialized bit patterns (Manchester and framing errors). In this test the receivers ability to detect these errors was also verified.

TEST 5: Protocol messages requiring the terminal to respond (SOLICIT_ENTRY and TOKEN) were sent. Appropriate responses (REQUEST ENTRY/STATUS_SUMMARY or TOKEN) were monitored from the terminal.

TEST 6: Messages that defined the terminals predecessor and successor in the topology map were sent (SET_PREDECESSOR, SET_SUCCESSOR, and SET_ TOPOLOGY). Registers in the terminal were monitored to verify that the appropriate registers had been updated.

TEST 7: Messages that loaded the reference clock timer (SET_CLOCK) were sent. The timer registers were monitored to verify that the time value had been updated.

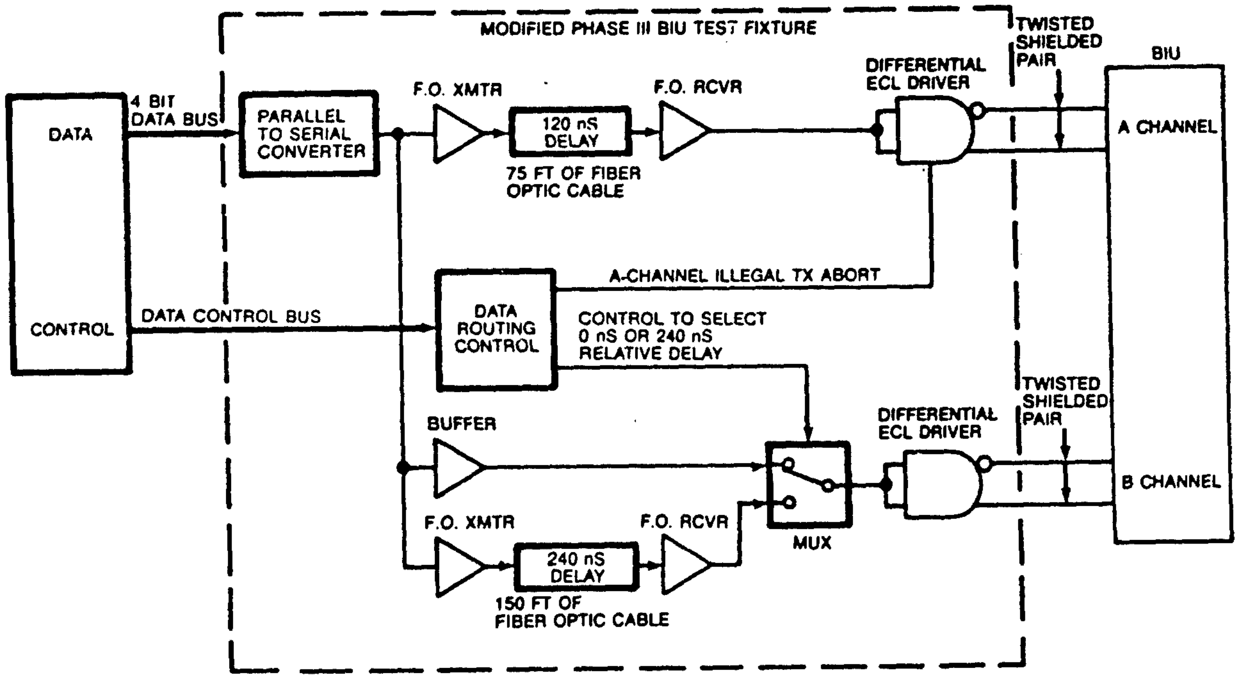

Figure 102. Brassboard Terminal Connected in the Data Generator Test Configuration

TEST 8: Data messages were sent to the terminal addressed using each of the prescribed addressing methods.

MH Card Testing

Hardware comprising the DPM function was tested using a Rockwell developed (not part of this program) AAMP test station. The test station was operated in a manner similar to that used to test most microprocessor based designs; the microprocessor chip was replaced with a specialized test pod which allowed direct interaction with the microprocessor while in operation. This allowed specific memory locations to be written, read, and monitored. It also allowed the test operator to set breakpoints and review operations prior to and after the breakpoint was encountered.

6.3.2 Network Management

Network management functions are those which establish and maintain the logical ring. They may be roughly organized into three groups: (1) steady state operation, (2) initialization and (3) error recovery.

Steady-State Operation

Steady-state operation of the network is shown in Figure 103. This shows three terminals exchanging tokens in a 3-node logical ring. You can see that the three distinctly different amplitudes allow identification of the node which generated each token. This mode of operation is initiated following startup of the logical ring by the RINGMASTER terminal, and continues until one of the terminals has a message to send or a failure occurs.

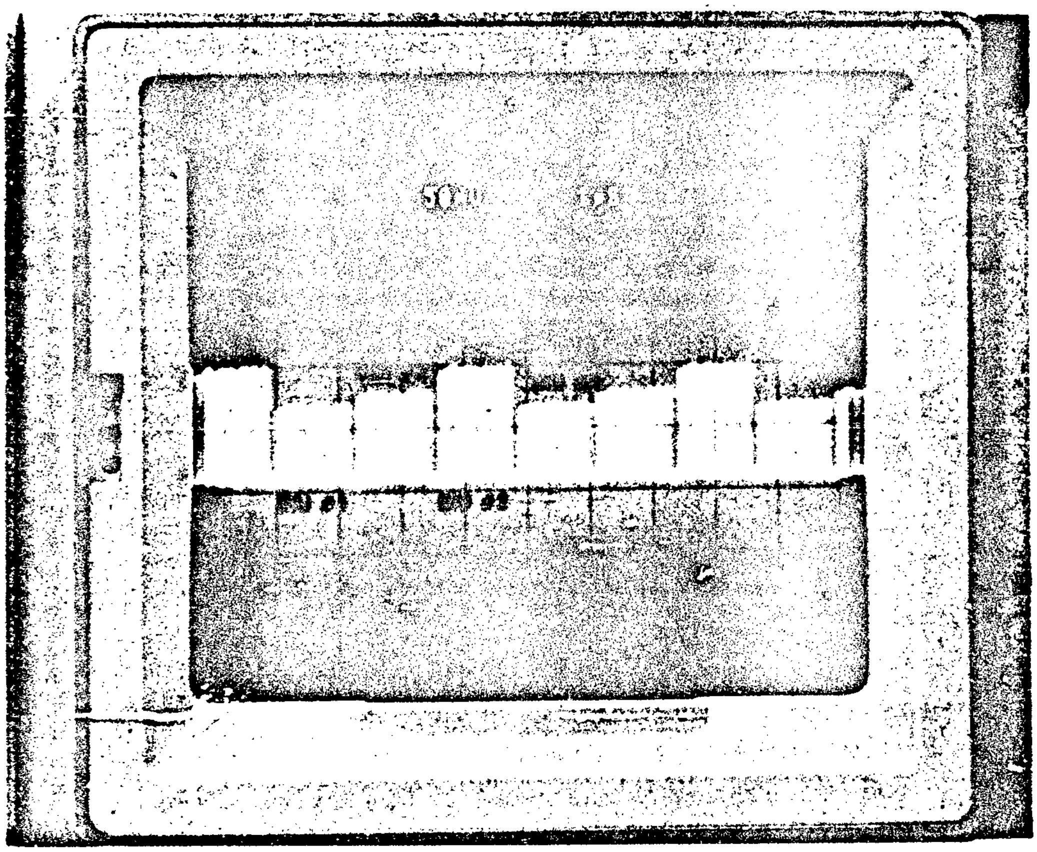

Figure 103. Three Node HSDB in Operation

Figure 104 shows an expanded view of a single token waveform. The token begins with a preamble (in this case 4 bits). Next comes the start delimiter, then frame control, destination address, source address, frame check sequence, and finally the end delimiter. This gives a total token time of 72 bits.

Figure 105 shows the critical intermessage gap timing of the breadboard terminal. The system specification requires a terminal response to a token to occur between 100 nS and 200 nS from when the token was received. The brassboard terminal response was 130 nS, worst case.

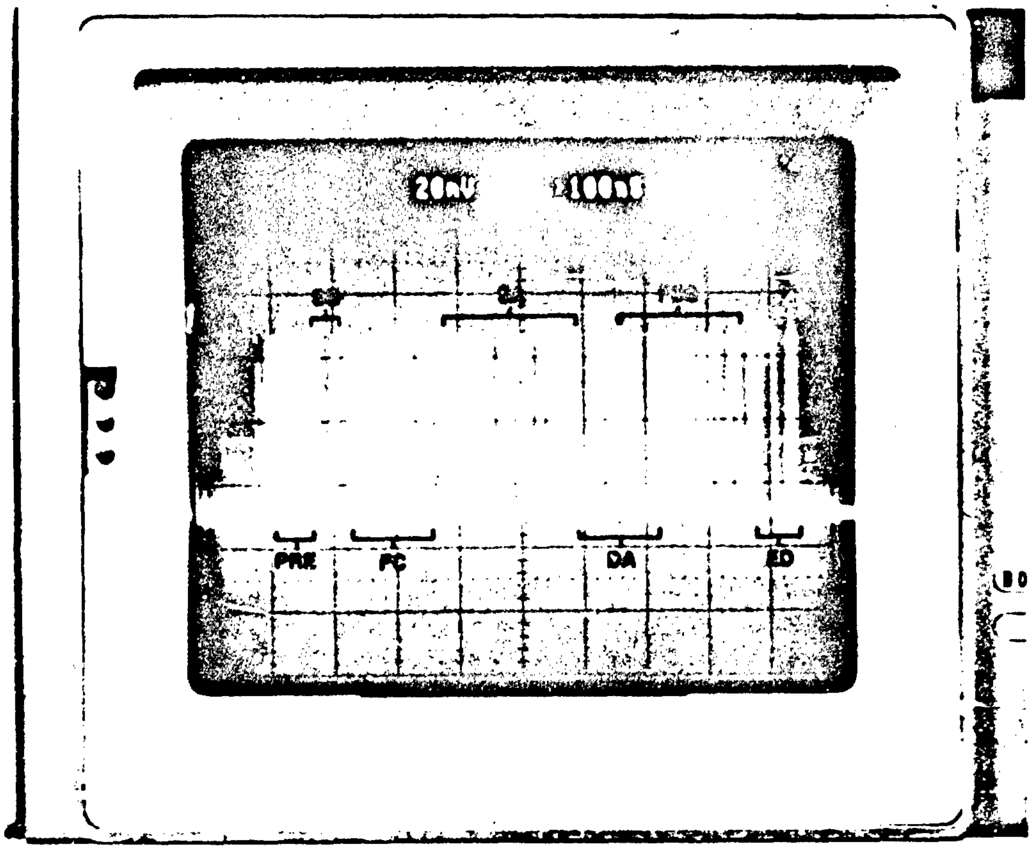

Figure 104. Expanded Token Packet Showing Construction

Initialization

Initialization refers to the process of organizing and starting steady state operation of the network from a non-operating configuration such as turn-on. The HSDB system specification describes two different algorithms for accomplishing this function, the collision vie process and the clear token vie process. The breadboard terminal implemented the collision vie process; the brassboard terminal implemented the clear token process.

Figure 106 shows the collision vie process. Upon startup or detection of a network activity error, each terminal transmits a special claim token message whose length is proportional to its physical address. Following transmission of its claim token, it listens for network activity. If it senses activity, it defers to the terminal with the higher address, if it hears a quiet network, it begins to poll other network addresses to establish which nodes are active. The first part of the waveform shows the collision as a high-level signal. First three transmitters are additive, then two and finally one. The one capturing the token is then seen to send a series of solicit entry messages, one to each address lower than its own. The figure shows the response on the part of the other two active nodes. Figure 107 shows this more clearly. The terminal polled responds to the solicit entry message from the RINGMASTER with a request entry and status message. The

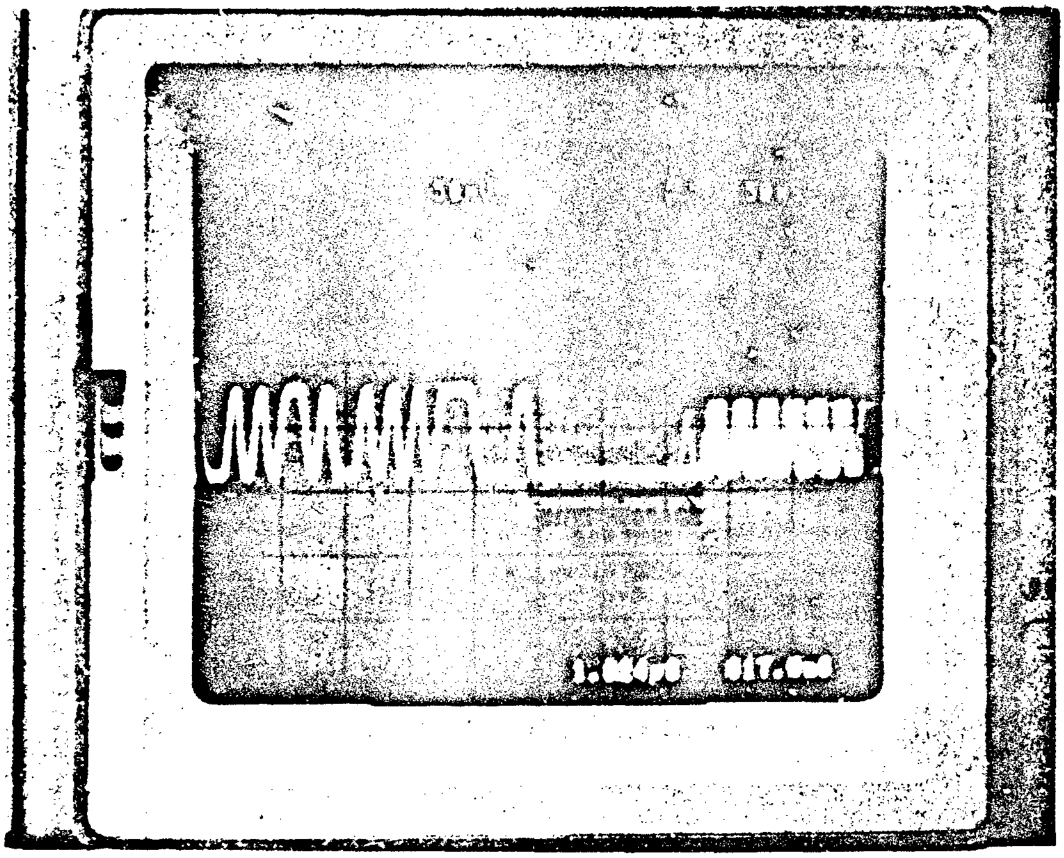

Figure 105. The Terminal Response Time Requirement has been Validated

RINGMASTER then sends SET_PREDECESSOR and SET_SUCCESSOR messages to splice the requesting node into the logical ring. The sequence terminates with the RINGMASTER broadcasting the topology map and then passing the token to initiate steady state operation. This sequence is shown on the far right of Figure 108.

Error Recovery

The need for and process of error recovery is illustrated in Figure 108. As shown, the terminal holding the token attempted to pass it to its successor. The successor would normally accept the token and begin transmitting within 200 nS. In this test case, the destination terminal does not respond to the first attempt at the token pass. This is shown by the large dead space in network activity. The holding terminal, after waiting a period defined by the response window timer, retries the token pass. In this (the second) try, the destination accepts the token and steady state operation ensues.

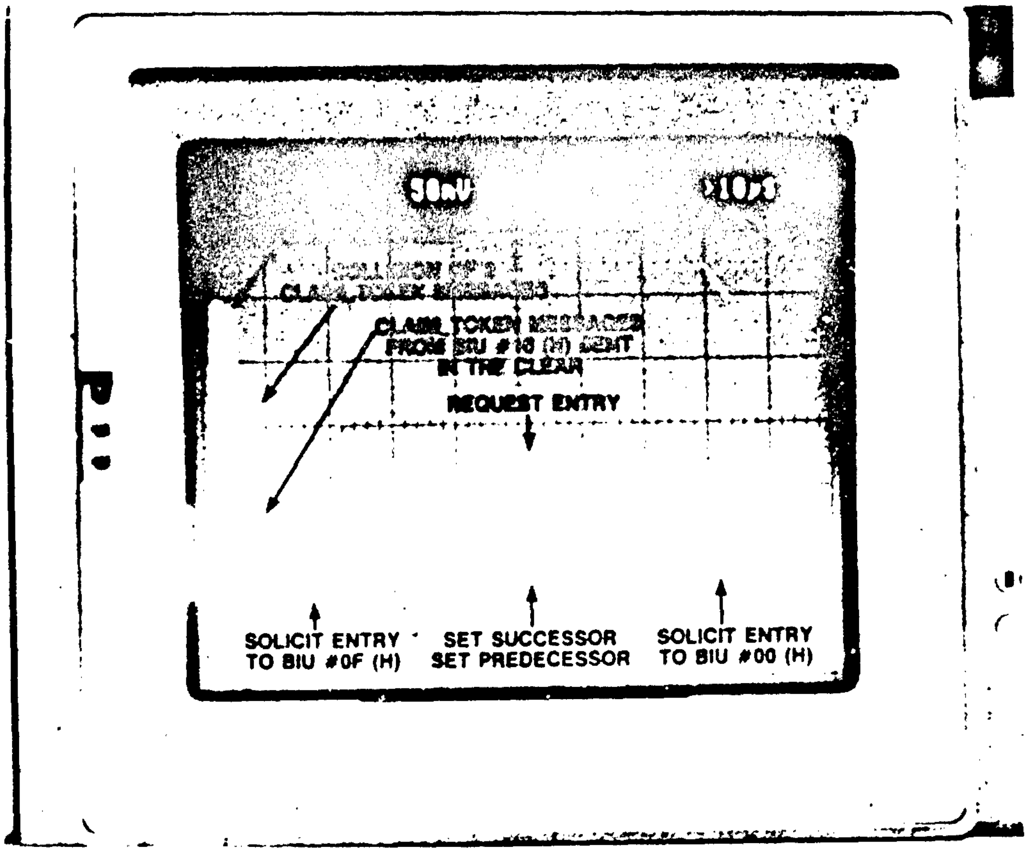

Figure 106. The Collision Vie Process has been Demonstrated

6.3.3 Service Functions

Service functions are attributes of the protocol which are not directly related to operation of the network. Service functions are described in the PAVE PILLAR HSDB system specification:

a. Reference clock
b. Statistics
c. Redundancy

These have all been validated as part of the brassboard development phase of Task IV.

Reference Clock

The reference clock timer is physically located in the IC gate array of the brassboard. It consists of 3 16-bit registers operating from a 1 μs clock. The registers may be set from the HSDB, read by the OC gate array (to generate a clock message for presentation to the network) set from the message handler, or read by the message handler.

Test/verification of this function was performed in conjunction with test of the E1 brassboard.

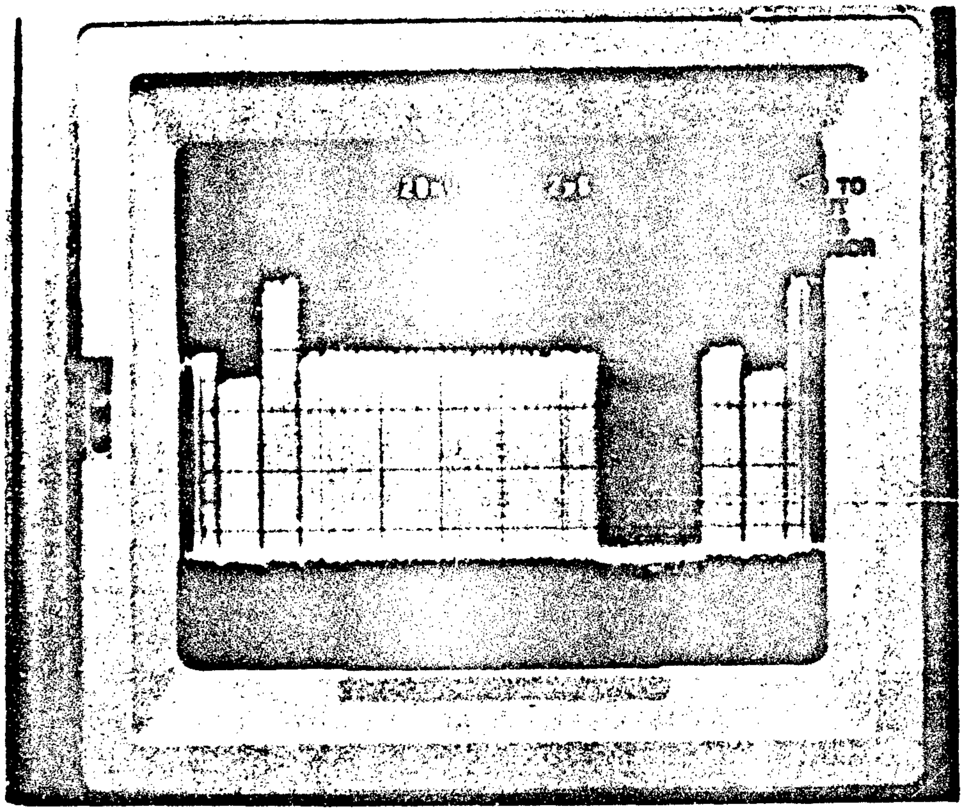

Figure 107. Demonstration of the Handshake Between the Ringmaster and a Terminal Requesting Network Entry

Statistics

The PAVE PILLAR HSDB system specification contains the requirement that a variety of terminal and network statistics be maintained by each terminal. These statistics may be retrieved either by the local user or by the user associated with any other HSDB node. The brassboard terminal design includes the entire set of 33 statistics functions. These are implemented using a variety of means, dedicated registers embedded in the gate arrays, general purpose hardware registers co-located with the message handler processor, and software registers accessible by the message handler processor.

Test/verification of this function was performed in conjunction with test of the E1 brassboard terminal.

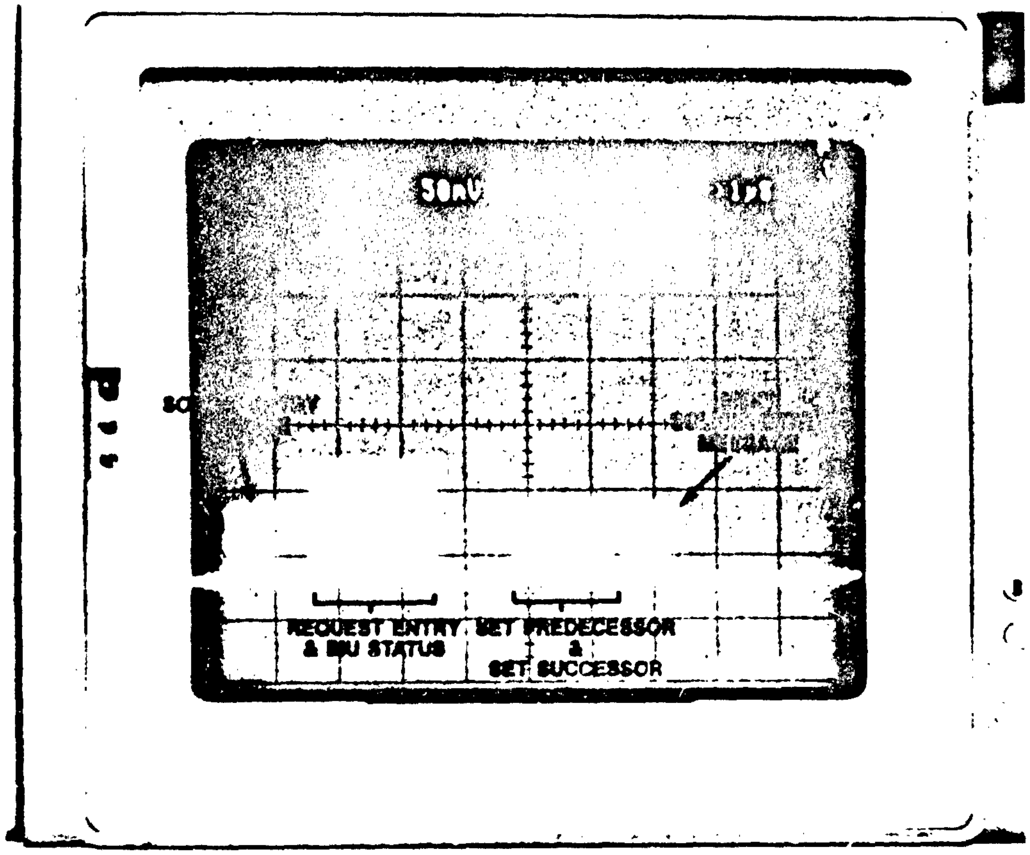

Figure 108. Demonstration of Dropped Token Recovery by Retry of the Token Pass

Redundancy

The standard architecture of a HSDB network is described as having dual path synchronous redundancy. This means that two completely independent paths exist through the network and that all messages are sent (and received) synchronously on both paths. The terminal at each node must contain logic which receives the packet from each network and recovers the first of the redundant messages to arrive without error.

This redundancy algorithm is included in the design of the brassboard terminals and was verified as part of engineering lab testing performed on the E1 brassboard terminal.

6.4 HSDB System Demonstration Equipment

The HSDB demonstration system shown in Figures 34, 63 and 109 is a complex design which represented a major part of Task IV development effort. The system contains items of commercial test equipment, government furnished equipment, special test equipment, software, and prototype HSDB terminals interconnected as shown in Figure 109. Each user processor consisted of a Sperry 1631 computer with a cathode ray terminal and a dual floppy disk drive as

peripherals. Special software was developed to simulate a typical user. This allowed messages to be sent to another user and received from another user, emulating activity patterns expected on an operational HSDB network. Each user processor was served by a dedicated prototype HSDB terminal. The terminals implemented the network protocol, formatted messages into packets and placed them on the network, and recovered messages from the network addressed to the local user. Terminals also provided access points to internal signals for connection with test equipment such as shown in the block diagram. This allows network operation to be accurately characterized. Two generations of the HSDB demonstration system were designed. The first, using breadboard hardware, implemented a demonstration protocol. This generation of the equipment was used to validate and demonstrate wire and fiber optic TRU technology in conjunction with the Task I and Task II ATRs.

The second generation of the system was designed around the breadboard terminals. These implemented Version 3.0 of the PAVE PILLAR HSDB protocol. This generation of the equipment was used to validate and demonstrate core functions of the PAVE PILLAR protocol including network initialization, token passing, message transmission and reception, and recovery from a dropped token.

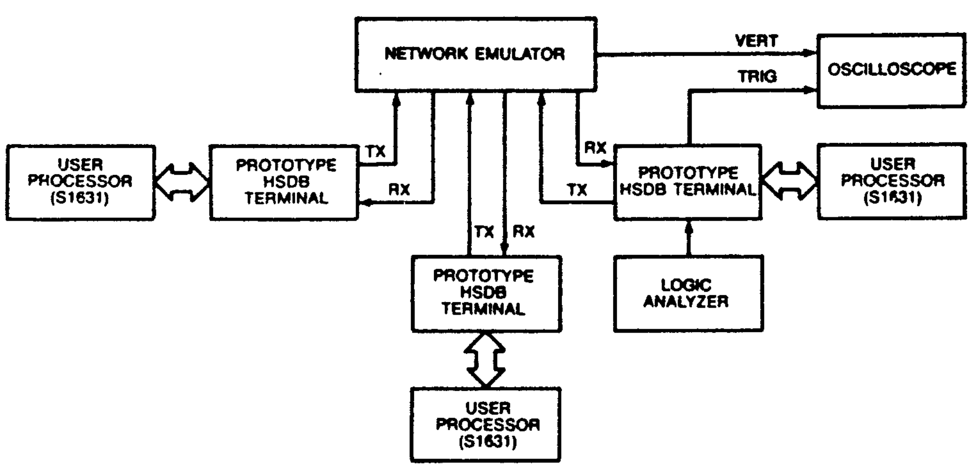

Figure 109. The System Demonstration Equipment Block Diagram

7.0 RESULTS, CONCLUSIONS, AND RECOMMENDATIONS

The HSB Technology Development Pr.)gram has been a success. The two major objectives of the program have been met.

1. Enabling technologies have been developed which show the practicality of operating a 50 Mbps data bus with 64 nodes using either coaxial cable or fiber optic interconnect.
2. An efficient and reliable protoc.l for the HSDB has been developed and demonstrated.

During the course of this program, the most sign.fi. .nt risks associated with the use of a local area network aboard an aircraft have been addresseo .nd solutions have been demonstrated. As a result, Rockwell feels that the state-of-the-art is such that a 64-node HSDB can be developed for next generation aircraft at minimal cost and risk.

Results

Several significant technological advances have resulted from this program. These advances are visible both in the direct results of this program and also in other related programs including ATF DEM/VAL hardware, LHX designs and SAE standards. Rockwell believes that this program has done much to lead HSDB technology from a point where it was considered to be very high risk, to a point where today there is little doubt that a fiber optic HSDB will form the backbone communication network for the next generation of military aircraft. Some of the more significant results are summarized below.

- A passive coaxial linear bi-directional coupler which allows reliable operation of a 64-node HSDB network aboard military aircraft has been developed, characterized, and demonstrated. Prior to this program such a design was widely considered to be impossible above 20 Mbps.
- Fiber optic transmitter and receiver designs which operate over the temperature range from -54 ^{o}C through +95 ^{o}C has been developed, characterized, and demonstrated. Prior to this program no receivers meeting the combination of sensitivity, dynamic range, data rate, error rate, and temperature range of operation had been produced. Also, no optical transmitter meeting the peak output power requirement over temperature had been produced.
- An efficient and reliable token passing protocol has been designed, characterized, and demonstrated using a test message data base derived from a survey

of user requirements. The protocol has been analyzed for latent defects using a formal lock-up analysis.

- The impact on network latency of a variety of protocol and network parameters has been determined using a computer hosted simulation tool. This allowed the latency control mechanism and the initialization/recovery mechanism of the protocol to be analyzed and optimized for the target application. Prior to this program no such formal analyses had been performed and, consequently, the quality of protocol designs was open to question.

Conclusions

As a result of performing the program, a number of conclusions have been reached. These are summarized here in order to guide planning for those contemplating related programs, especially development of production designs.

- Development of a coaxial cable interconnected network is practical. At the inception of the program it was considered a high risk technology. This program produced the design for a passive coupler for a tapped linear bus topology which is reliable and reproducible.
- Development of a fiber optic interconnected network is practical although not as straight forward as is development of a coaxial network. This program produced a design for a receiver exhibiting state-of-the-art sensitivity and dynamic range performance. The sensitivity goal established early during the design was never quite achieved. Fortunately, the state of the art of LED sources improved faster than expected allowing the planned power budget for the network to be achieved. Production HSDB designs will need to carefully balance requirements for receiver sensitivity, transmitter power output, and network topology/loss to arrive at a reliable operating point. In many applications it may not be possible to use a passive interconnect if more than 12 nodes are required.
- The PAVE PILLAR protocol is reliable, efficient, and well behaved under all operating conditions. The priority mechanism allows latency control for critical messages while causing a minimum of degradation to non-priority messages. Most applications will not require the use of the latency control mechanism, the network operates so well to above 30 Mbps average throughput that little enhancement will be realized from the use of priorities.

Recommendations

From experience gained from this program, Rockwell recommends that the Air Force initiate two additional activities. These address areas which are not covered within the scope of the present contract but which we consider to be highly desirable in order to better prepare for production applications.

- The design, development, and demonstration of a flightworthy terminal meeting the PAVE PILLAR specification should be completed. This would allow comparison with other candidate HSDB designs such as those used for ATF DEM/VAL. The first generation of chips designed under this contract have shown that this can be accomplished at relatively low cost and risk.
- The development of a fiber optic linear tapped bus topology should be initiated. This program has proven the feasibility of fiber optic TRU designs for use abroad aircraft. The interconnect used, however, is not adequate for the majority of planned aircraft installations. This program demonstrated a power budget of approximately 30 dB; a high performance tactical aircraft will require somewhere in the order of 50 dB. While it seems straightforward to simply replace the passive star coupler with an active star coupler to provide the needed power budget, this ignores the characteristics of the intended installation. Running a large quantity of fibers through the airframe is undesirable and will reduce the reliability of the network. A linear topology similar to that of MIL-STD-1553B systems is much better from the perspective of aircraft design, and should be developed. Rockwell believes this to be a realizable objective.

REFERENCES

1. MIL-STD-1553B, "Military Standard, Aircraft Internal Time Division Command/Response Multiplex Data Bus", 21 September, 1978

2. IEEE 802.3-1985, "Carrier Sense, Multiple Access With Collision Detection", 23 June, 1983

3. IEEE 802.4-1985, "Token Passing Bus Access Method", 13 June, 1984

4. SAE AE-9B/L, "Linear Token-Passing Multiplex Bus Standard", Draft C

5. Report on High Speed Data Bus Application And Requirements Task (HART) for the SAE AE-9B/L High Speed Data Bus

6. "PCM And Digital Transmission Systems", McGraw-Hill Company, 1982

7. MIL-E-5400T, "Military Specification, Electronic Equipment, Aerospace, General Specification For", 16 November, 1979

8. K. Ogawa and E. L. Chinnock, "GaAs Transimpedance Front-end Design For A Wideband Optical Receiver", Electronic Letters, 1979

9. PAVE PILLAR HSDB System Specification USAF Contract F33615-83-C-1036, CDRL No. 10-3 (Version 3 included in "Architecture Specification for PAVE PILLAR Avionics", Report No. AFWAL-TR-87-1114, January 1987, DTIC AD-A188-722)

BIBLIOGRAPHY

Ferris, K., High Speed Fiber Optic Data Bus For Avionics Integration, paper presented at IEEE National Aerospace and Electronics Conference (NAECON) 85, Dayton, Ohio, May 1985

Goldman, P.C. and Ludvigson, M.T., High Speed Bus For PAVE PILLAR Applications, paper presented at SAE Aerospace Avionics Equipment and Integration Committee (AE-9) Technical Symposium, Phoenix, Arizona, April 1986

Ludvigson, M.T. and Milton, K.L., A High Speed Bus For PAVE PILLAR Applications, paper presented at SAE Aerospace Congress and Exposition, Long Beach, California, October 1984

Ludvigson, M.T. and Milton, K.L., A High Speed Bus for PAVE PILLAR Applications, paper presented at NAECON 85 Conference, Dayton, Ohio, May 1985

Ludvigson, M.T. and Milton, K.L., A High Speed Bus For PAVE PILLAR Applications, paper presented at AIAA Computers in Aerospace Conference, Long Beach, California, October 1985

Ludvigson, M.T., On High Speed Data Bus Performance Analysis, and, Token Rotation Timer Implementation, papers presented at SAE Avionics Systems Division (ASD) Technical Symposium, Dallas, Texas, November 1987

Ludvigson, M.T., On High Speed Data Bus Performance Analysis, (updated version) paper presented at NAECON 88 Conference, Dayton, Ohio, May 1988

Ludvigson, M.T., Modrow, M.B., and Goldman, P.C., High Speed Bus Technology Development Program, paper presented at AIAA/IEEE Digital Avionics Systems Conference (DASC), San Jose, California, October, 1988

Modrow, M.B., Ludvigson, M.T. and Goldman, P.C., A High Speed Data Bus Technology Development Program, paper presented at Armed Forces Communications Electronics Association (AFCEA), Norfolk, Virginia, November 1987

Rockwell and USAF, Implementation Considerations For the High Speed Data Bus, paper distributed to SAE AE-9B White Paper Evaluation Board, March 1984

Modrow, M.B., High Speed Data Bus For the Air Force, AVIONICS Magazine, March 1989

Ludvigson, M.T., Bursty Traffic Bus Performance Analysis, paper presented at SAE ASD Technical Symposium, Seattle, Washington, April 1989

Ludvigson, M.T. and Goldman, P.C., A High Speed Bus For Vehicle And Mission Control, paper presented at the Association for Unmanned Vehicle Systems Spring Symposium, Johns Hopkins University, Laurel, Maryland, May 1986

*U.S.Government Printing Office: 1990—748-056/24396